Non-Asymptotic Delay Bounds for Multi-Server Systems with Synchronization Constraints

Markus Fidler, Brenton Walker, and Yuming Jiang

arXiv:1610.06309v1 [cs.PF] 20 Oct 2016

***Abstract*—Multi-server systems have received increasing attention with important implementations such as Google MapReduce, Hadoop, and Spark. Common to these systems are a fork operation, where jobs are first divided into tasks that are processed in parallel, and a later join operation, where completed tasks wait until the results of all tasks of a job can be combined and the job leaves the system. The synchronization constraint of the join operation makes the analysis of fork-join systems challenging and few explicit results are known. In this work, we model fork-join systems using a max-plus server model that enables us to derive statistical bounds on waiting and sojourn times for general arrival and service time processes. We contribute end-to-end delay bounds for multi-stage fork-join networks that grow in $\mathcal{O}(h \ln k)$ for h fork-join stages, each with k parallel servers. We perform a detailed comparison of different multi-server configurations and highlight their pros and cons. We also include an analysis of single-queue fork-join systems that are non-idling and achieve a fundamental performance gain, and compare these results to both simulation and a live Spark system.**

I. Introduction

Fork-join systems are an essential model of parallel data processing, e.g., Google MapReduce [2], Hadoop, or Spark [3], where jobs are divided into k tasks (fork) that are processed in parallel by k servers. Once all tasks of a job are completed, the results are combined (join) and the job leaves the system. Fig. 1 illustrates an example. Multi-stage fork-join networks comprise several fork-join systems in tandem, where all tasks of a job have to be completed at the current stage before the job is handed over to the next stage. The difficulty in analyzing such systems is due to a) the statistical dependence of the workload on the parallel servers that is due to the common arrival process [4], [5], and b) the synchronization required by the join operation [4], [6].

Significant research has been performed to analyze the performance of fork-join systems. However, exact results are known only for few specific systems, such as two parallel M|M|1 queues [7], [8]. For more complex systems, approximation techniques, e.g., [5], [8]–[14], and bounds, using stochastic orderings [4], martingales [15], or stochastic burstiness constraints [16], have been explored. Given the difficulties posed by single-stage fork-join systems, few works consider multi-stage networks. A notable exception is [13] where an approximation for closed fork-join networks is developed.

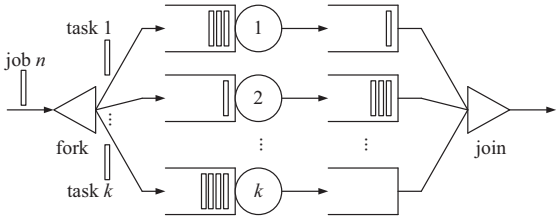

Fig. 1. Fork-join system. Each job is composed of k tasks with individual service requirements, that are distributed to k servers (fork). Once all tasks of a job are completed, the job leaves the system (join), i.e., the tasks of a job wait at the join step until all tasks of the job are completed.

Related synchronization problems also occur in the case of load balancing using parallel servers and in the case of multi-path routing of packet data streams [17] using multi-path protocols [15]. The tail behavior of delays in multi-path routing is investigated in [17] as well as in [18], [19] where large deviation results of resequencing delays for parallel M|M|1 queues are derived.

Split-merge systems are a variant of fork-join systems with a stricter synchronization constraint: all tasks of a job have to start execution simultaneously. In contrast, in a fork-join system, the start times of tasks are not synchronized. Split-merge systems are solvable to some extent as they can be expressed as a single server queue where the service process is governed by the service time of the maximal task of each job [9], [15], [20], [21].

Most closely related to this work are three recent papers [15], [16], [22] that employ similar methods. The work [16] considers single-stage fork-join systems with load balancing, general arrivals of the type defined in [23], and deterministic service. A service curve characterization of fork-join systems is provided and statistical delay bounds are presented. The paper [15] contributes delay bounds for single-stage fork-join systems with renewal as well as Markov-modulated inter-arrival times and independent and identically distributed (iid) service times. The authors prove that delays for fork-join systems grow as $\mathcal{O}(\ln k)$ for k parallel servers, as also found in [4]. Split-merge systems are shown to have inferior performance [15], where the stability region of k parallel M|M|1 queues decreases with $\ln k$. The work also includes a first application to multi-path routing, assuming a generic window-based protocol that operates on batches of packets. The authors conclude that multi-path routing is only beneficial in the case of two parallel paths and mod-

This work was supported in part by the European Research Council (ERC) under Starting Grant UnIQue (StG 306644). This manuscript is a revised and extended version of the paper [1] that appeared in the IEEE Infocom 2016 proceedings.

M. Fidler and B. Walker are with the Institute of Communications Technology, Leibniz Universität Hannover, Germany, (E-mail: markus.fidler@ikt.uni-hannover.de and brenton.walker@ikt.uni-hannover.de). Y. Jiang is with the Department of Telematics, NTNU Trondheim, Norway (E-mail: jiang@item.ntnu.no).

erate to high utilization. Otherwise, resequencing delays are found to dominate. In [22], the authors evaluate different task assignment policies for parallel server systems with task replication, considering the effects of correlated replicas. Task replication relates to the more general concept of (k, l) fork-join systems [1], [21], where a job is considered completed once l out of k tasks have finished service.

While [15] focuses on split-merge vs. fork-join systems with iid service times, we also consider the case of non-iid service, where we are able to generalize important results, such as the growth of delays in $\mathcal{O}(\ln k)$ for fork-join systems with k parallel servers. Furthermore, we show that fork-join systems can be formulated as a server under the max-plus algebra [24]. This essential lemma enables the analysis of multi-stage fork-join networks. For h statistically independent fork-join stages each with k parallel servers, we prove that the growth of delays is in $\mathcal{O}(h \ln k)$. The result compares to a scaling in $\mathcal{O}(h \ln(hk))$ that we obtained previously in [1] without assuming independence of the stages.

We perform a detailed evaluation of different multi-server configurations which reveals that fork-join systems mostly but not universally outperform classical multi-server systems. Beyond [1], [15], [16], [22], we also include single-queue multi-server as well as single-queue fork-join systems. In contrast to the standard fork-join model, where each of the servers has an individual queue, single-queue systems are non-idling in the sense that queueing can occur only if all servers are busy. Our evaluation reveals a fundamental performance gain of non-idling single-queue systems. We include reference results, mostly obtained by simulation as well measurements from a live Spark cluster, that verify the tightness of our performance bounds.

The remainder of this paper is structured as follows. In Sec. II, we formulate basic models of G|G|1 as well as GI|GI|1 fork-join systems in max-plus system theory. Multi-stage fork-join networks are considered in Sec. III. We compare fork-join systems with classical multi-server systems with thinning and optional resequencing in Sec. IV. In Sec. V, we analyze non-idling single-queue implementations of multi-server and fork-join systems, respectively. Sec. VI presents brief conclusions. Extensive proofs and a detailed description of the simulation and Spark experiments are in the appendix.

II. Basic Fork-Join Systems

In this section, we derive a set of results for basic fork-join systems in max-plus system theory [24]–[28]. Max-plus system theory is a branch of the deterministic [24], [29], [30], respectively, stochastic network calculus [24], [31]–[37]. In comparison to [15], which is focused entirely on waiting and sojourn times of specific systems, the more general max-plus approach enables us to construct multi-stage fork-join networks as well as more advanced fork-join systems. Further, we generalize central results from [15] considering general arrival and service processes. Throughout this work, we consider only the case of homogeneous servers, i.e., all servers have identical service time distribution. Heterogeneous servers can be dealt with in the same way by a notational extension. We show results for heterogeneous servers and load balancing in [1].

A. Notation and Queueing Model

We label jobs in the order of arrival by $n \geq 1$ and let $A(n)$ denote the time of arrival of job n. It follows for $n \geq m \geq 1$ that $A(n) \geq A(m) \geq 0$. For notational convenience, we define $A(0) = 0$. Further, we let $A(m, n) = A(n) - A(m)$ be the time between the arrival of job m and job n for $n \geq m \geq 1$. Hence, $A(n, n+1)$ is the inter-arrival time between job n and job $n+1$ for $n \geq 1$. Similarly, $D(n)$ denotes departure times. To model systems, we adapt the definition of g-server from [24, Def. 6.3.1] using a notion of service process $S(m, n)$.

Definition 1 (Max-plus server)**.** *A system with arrivals $A(n)$ and departures $D(n)$ is an $S(m, n)$ server under the max-plus algebra if it holds for all $n \geq 1$ that*

$$D(n) \leq \max_{m \in [1,n]} \{A(m) + S(m, n)\}.$$

It is an exact $S(m, n)$ server if it holds for all $n \geq 1$ that

$$D(n) = \max_{m \in [1,n]} \{A(m) + S(m, n)\}.$$

The following Lem. 1 shows that the general class of work-conserving systems satisfy the definition of exact server. We use $V(n)$ to denote the time at which job n starts service.

Lemma 1 (Work-conserving system)**.** *Consider a lossless, work-conserving, first-in first-out system and let $L(n)$ denote the service time of job n, where $n \geq 1$. Define for $n \geq m \geq 1$*

$$S(m, n) = \sum_{\nu=m}^{n} L(\nu).$$

The system is an exact $S(m, n)$ server.

Proof: Since the system is lossless, work-conserving, and serves jobs in first-in first-out order, job $n \geq 2$ starts service at

$$V(n) = \max\{A(n), V(n-1) + L(n-1)\}, \tag{1}$$

and job 1 at $V(1) = A(1)$. By recursive insertion of (1) we have

$$V(n) = \max_{m \in [1,n]} \left\{ A(m) + \sum_{\nu=m}^{n-1} L(\nu) \right\}, \tag{2}$$

for $n \geq 1$. Since $D(n) = V(n) + L(n)$, it follows with (2) that $D(n) = \max_{m \in [1,n]}\{A(m) + \sum_{\nu=m}^{n} L(\nu)\}$, which proves that the work-conserving system is an exact max-plus server. ∎

For the sojourn time of job $n \geq 1$, defined as $T(n) = D(n) - A(n)$, it follows by insertion of Def. 1 that

$$T(n) = \max_{m \in [1,n]} \{S(m, n) - A(m, n)\}. \tag{3}$$

The waiting time of job $n \geq 1$ is $W(n) = V(n) - A(n)$. As in the case of work-conserving systems in Lem. 1, $V(n) = \max\{A(n), D(n-1)\}$, so we have $W(n) = [D(n-1) -$

$A(n)]^+$, where $[X]^+ = \max\{X, 0\}$ is the non-negative part and $D(0) = 0$ by definition. With Def. 1, it holds that

$$W(n) = \left[\sup_{m \in [1, n-1]} \{S(m, n-1) - A(m, n)\} \right]^+. \quad (4)$$

Here, we use the supremum since for $n = 1$ (4) evaluates to an empty set. For non-negative real numbers the sup of an empty set is zero. While we used the definition of an exact server to derive an expression for the sojourn and waiting times, we note that the upper bound specified by the definition of a server is usually sufficient, as it provides upper bounds of sojourn and waiting times.

B. Statistical Performance Bounds

Next, we derive statistical performance bounds for servers as defined above. Throughout the paper, we generally assume that the arrival and service processes are independent of each other. Considering general arrival and service processes, the server is a G|G|1 queue. The results enable us to generalize recent findings obtained for iid service times, i.e., for a GI service model, in [15].

We consider arrival and service processes that belong to the broad class of (σ, ρ)-constrained processes [24], that are characterized by affine bounding functions of the moment generating function (MGF). The MGF of a random variable X is defined as $\mathsf{M}_X(\theta) = \mathsf{E}\left[e^{\theta X}\right]$ where θ is a free parameter. The following definition adapts [24] to max-plus systems.

Definition 2. *An arrival process is (σ_A, ρ_A)-lower constrained if for all $n \geq m \geq 1$ and $\theta > 0$ it holds that*

$$\mathsf{E}\left[e^{-\theta A(m,n)}\right] \leq e^{-\theta(\rho_A(-\theta)(n-m)-\sigma_A(-\theta))}.$$

Similarly, a service process is (σ_S, ρ_S)-upper constrained if for all $n \geq m \geq 1$ and $\theta > 0$ it holds that

$$\mathsf{E}\left[e^{\theta S(m,n)}\right] \leq e^{\theta(\rho_S(\theta)(n-m+1)+\sigma_S(\theta))}.$$

Considering the service times of jobs as in Lem. 1, we also apply Def. 2 to characterize the cumulative service process $L(m, n) = \sum_{\nu=m}^{n} L(\nu)$ by (σ_L, ρ_L).

In the special case of GI arrival processes, $A(m, n) = \sum_{\nu=m}^{n-1} A(\nu, \nu+1)$ has iid inter-arrival times $A(\nu, \nu+1)$. It follows that $\mathsf{E}\left[e^{-\theta A(\nu,\nu+1)}\right] = \mathsf{E}\left[e^{-\theta A(1,2)}\right]$ for $\nu \geq 1$. Next, we use that the MGF of a sum of independent random variables is the product of their individual MGFs, i.e., $\mathsf{M}_{X+Y}(\theta) = \mathsf{M}_X(\theta)\mathsf{M}_Y(\theta)$ to derive minimal traffic parameters from Def. 2 as $\sigma_A(-\theta) = 0$ and

$$\rho_A(-\theta) = -\frac{1}{\theta} \ln \mathsf{E}\left[e^{-\theta A(1,2)}\right]. \quad (5)$$

Similarly for GI service processes, $S(m, n) = \sum_{\nu=m}^{n} S(\nu)$ is composed of iid service increments $S(\nu)$ that have minimal parameters $\sigma_S(\theta) = 0$ and

$$\rho_S(\theta) = \frac{1}{\theta} \ln \mathsf{E}\left[e^{\theta S(1)}\right]. \quad (6)$$

Parameter $\rho_A(-\theta)$ decreases with $\theta > 0$ from the mean to the minimum inter-arrival time and $\rho_S(\theta)$ increases with $\theta > 0$ from the mean to the maximum service time.

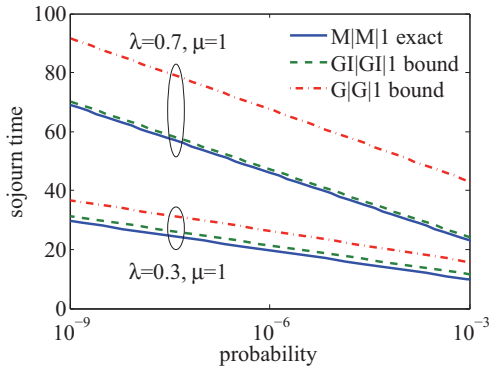

Fig. 2. M|M|1 queue. The bounds show the correct exponential tail decay.

Theorem 1 (Statistical performance bounds)**.** *Consider a server as in Def. 1, with arrival and service parameters $(\sigma_A(-\theta), \rho_A(-\theta))$ and $(\sigma_S(\theta), \rho_S(\theta))$ as specified by Def. 2. For $n \geq 1$, the sojourn time $T(n) = D(n) - A(n)$ satisfies*

$$\mathsf{P}[T(n) > \tau] \leq \alpha e^{\theta \rho_S(\theta)} e^{-\theta\tau},$$

and the waiting time $W(n) = [D(n-1) - A(n)]^+$ satisfies

$$\mathsf{P}[W(n) > \tau] \leq \alpha e^{-\theta\tau}.$$

In the case of G|G arrival and service processes, the free parameter $\theta > 0$ has to satisfy $\rho_S(\theta) < \rho_A(-\theta)$ and

$$\alpha = \frac{e^{\theta(\sigma_A(-\theta)+\sigma_S(\theta))}}{1 - e^{-\theta(\rho_A(-\theta)-\rho_S(\theta))}}.$$

In the special case of GI|GI arrival and service processes, $\theta > 0$ has to satisfy $\rho_S(\theta) \leq \rho_A(-\theta)$ and $\alpha = 1$.

The proof is provided in the appendix. For the special case of GI|GI arrival and service processes, Th. 1 recovers the classical bound for the waiting time of GI|GI|1 queues [38] in the max-plus system theory. Like [38], the proof uses Doob's martingale inequality [39]. The proof for the G|G arrival and service processes adapts the approach from [24], [33] to max-plus systems. The important property of the G|G result is that it differs only by a constant factor α from the GI|GI result and otherwise recovers the characteristic exponential tail decay $e^{-\theta\tau}$ with the same maximal decay rate θ.

M|M|1 Queue: For evaluation of the accuracy of the bounds in Th. 1, we consider the basic case of an M|M|1 queue, where exact results are available for comparison. Given iid exponential inter-arrival and service times with parameters λ and μ, respectively, (5) and (6) evaluate to

$$\rho_A(-\theta) = -\frac{1}{\theta} \ln\left(\frac{\lambda}{\lambda+\theta}\right), \quad (7)$$

for $\theta > 0$, and

$$\rho_S(\theta) = \frac{1}{\theta} \ln\left(\frac{\mu}{\mu-\theta}\right), \quad (8)$$

for $\theta \in (0, \mu)$. From the condition $\rho_S(\theta) \leq \rho_A(-\theta)$ it follows that $\theta \leq \mu - \lambda$ under the stability condition $\lambda < \mu$. By the choice of the maximal $\theta = \mu - \lambda$ we have from Th. (1) that

$$\mathsf{P}[T(n) > \tau] \leq \frac{\mu}{\lambda} e^{-(\mu-\lambda)\tau}. \quad (9)$$

Compared to the exact distribution of the sojourn time of the M|M|1 queue, that is $\mathsf{P}\left[T(n) > \tau\right] = e^{-(\mu-\lambda)\tau}$ see, e.g., [40], (9) has the same tail decay and differs only by the pre-factor μ/λ. Obviously, the bound becomes tighter if the utilization is high, in which case μ/λ approaches one.

In Fig. 2, we illustrate the bounds from Th. 1 compared to the exact M|M|1 result. Clearly, the curves show the same tail decay, where the GI|GI|1 bound provides better numerical accuracy compared to the G|G|1 bound that does not use independence of the increment processes and hence has parameter $\alpha > 1$. In the case of the G|G|1 bound, the parameter θ is optimized numerically to obtain the smallest delay bound.

For numerical evaluation, we will frequently use M|M arrival and service processes as specified by (7) and (8). We note that Th. 1 provides results for G|G arrival and service processes by substitution of the MGFs of the respective processes into Def. 2.

C. Fork-Join Systems

In a fork-join system, each job $n \geq 1$ is composed of k tasks with service times $Q_i(n)$ for $i \in [1,k]$, i.e., the service requirements of the tasks may differ from each other and may or may not be independent. The tasks are distributed to k parallel servers (fork) and once all tasks of a job are served, the job leaves the system (join), see Fig. 1. The parallel servers are not synchronized; i.e., server i starts serving task i of job $n+1$ (assuming it is already in the system), once it finishes serving task i of job n, which departs from server i at $D_i(n)$. Job n has finished service once all of its tasks $i \in [1,k]$ have finished service. The following lemma shows that fork-join systems are servers under the max-plus algebra. After estimating the MGF of the respective service process, performance bounds are obtained.

Lemma 2 (Fork-join system)**.** *Consider a fork-join system with k parallel servers as in Lem. 1. Let $Q_i(n)$ denote the service time of task i of job n where $i \in [1,k]$ and $n \geq 1$. Define for $n \geq m \geq 1$*

$$S(m,n) = \max_{i\in[1,k]}\left\{\sum_{\nu=m}^{n} Q_i(\nu)\right\}.$$

The system is an exact $S(m,n)$ server.

Proof: Since a job departs from the system once all of its tasks $i \in [1,k]$ are completed, we have for $n \geq 1$ that

$$D(n) = \max_{i\in[1,k]}\{D_i(n)\}. \tag{10}$$

By insertion of Def. 1 for each of the servers $i \in [1,k]$, it follows that

$$D(n) = \max_{i\in[1,k]}\left\{\max_{m\in[1,n]}\{A(m) + S_i(m,n)\}\right\}.$$

After reordering the maxima

$$D(n) = \max_{m\in[1,n]}\left\{A(m) + \max_{i\in[1,k]}\{S_i(m,n)\}\right\},$$

we conclude that the fork-join system is an exact $S(m,n) = \max_{i\in[1,k]}\{S_i(m,n)\}$ server. In the last step, we invoke Lem. 1 with $Q_i(n)$ for each of the servers $i \in [1,k]$. ∎

Next, we estimate the MGF of the service process $S(m,n)$ from Lem 2 for $n \geq m \geq 1$ by

$$\mathsf{E}\left[e^{\theta S(m,n)}\right] \leq \sum_{i=1}^{k} \mathsf{E}\left[e^{\theta\sum_{\nu=m}^{n} Q_i(\nu)}\right].$$

Assuming homogeneous tasks with parameters $(\sigma_Q(\theta), \rho_Q(\theta))$ for $i \in [1,k]$, it follows by insertion of Def. 2 that

$$\mathsf{E}\left[e^{\theta S(m,n)}\right] \leq k e^{\theta(\sigma_Q(\theta)+\rho_Q(\theta)(n-m+1))}.$$

This shows that the service process of the fork-join system has parameters

$$\sigma_S(\theta) = \sigma_Q(\theta) + \frac{\ln k}{\theta}, \tag{11}$$

and

$$\rho_S(\theta) = \rho_Q(\theta). \tag{12}$$

Performance bounds follow as a corollary of Th. 1.

Corollary 1 (Fork-join system)**.** *Consider a fork-join system as in Lem. 2, with arrival and service parameters $(\sigma_A(-\theta), \rho_A(-\theta))$ and $(\sigma_Q(\theta), \rho_Q(\theta))$ as specified by Def. 2. For $n \geq 1$, the sojourn time satisfies*

$$\mathsf{P}[T(n) > \tau] \leq k\alpha e^{\theta\rho_Q(\theta)} e^{-\theta\tau},$$

and the waiting time of the task that starts service last

$$\mathsf{P}\left[W(n) > \tau\right] \leq k\alpha e^{-\theta\tau}.$$

In the case of G|G arrival and service processes, the free parameter $\theta > 0$ has to satisfy $\rho_Q(\theta) < \rho_A(-\theta)$ and

$$\alpha = \frac{e^{\theta(\sigma_A(-\theta)+\sigma_Q(\theta))}}{1 - e^{-\theta(\rho_A(-\theta)-\rho_Q(\theta))}}.$$

In the special case of GI|GI arrival and service processes, $\theta > 0$ has to satisfy $\rho_Q(\theta) \leq \rho_A(-\theta)$ and $\alpha = 1$.

Proof: For G|G arrival and service processes, Cor. 1 is obtained directly by insertion of (11) and (12) into Th. 1. We note that the waiting time of a job is defined to be that of its task that starts service last. This follows by insertion of (10) into the definition of waiting time $W(n) = [D(n-1) - A(n)]^+$.

As the increment process of $S(m,n)$ in Lem 2 is non-trivial, we pursue a different approach[1] to show the result for the special case of GI|GI arrival and service processes. With (10), we derive the sojourn time $T(n) = D(n) - A(n)$ of job n as $T(n) = \max_{i\in[1,k]}\{D_i(n) - A(n)\}$ for $n \geq 1$. Hence, the sojourn time of job n is expressed as a maximum $T(n) = \max_{i\in[1,k]}\{T_i(n)\}$ of the sojourn times $T_i(n) = D_i(n) - A(n)$ of the individual tasks $i \in [1,k]$ of job n. Since the individual servers satisfy Lem. 1, we can invoke Th. 1 for each of the servers $i \in [1,k]$ and use the union bound to obtain the result

[1]The approach applies also in the case of G|G arrival and service processes. We showed the alternative approach via (11) and (12) nevertheless, as it extends to multi-stage fork-join networks, see Sec. III.

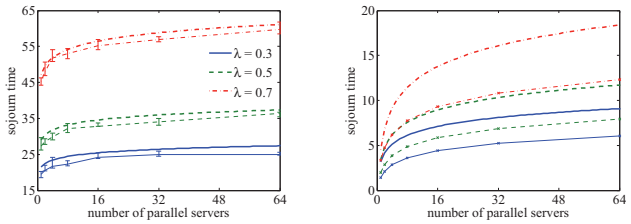

(a) $(1-\varepsilon)$-quantile (b) expected value

Fig. 3. Fork-join system. Analytical bounds (thick lines) and simulation results (thin lines). Sojourn times grow with $\ln k$ for k servers.

of Cor. 1. The waiting time $\max_{i\in[1,k]}\{W_i(n)\}$ where $W_i(n)$ as in (4) can be derived in the same way. ∎

We note that Cor. 1 does not make an assumption of independence regarding the parallel servers. Indeed, independence cannot be assumed as the waiting and sojourn times of the individual servers depend on the same arrival process [4], [5].

To investigate the scaling of fork-join systems with k parallel servers, we first note that the stability condition $\rho_Q(\theta) < \rho_A(-\theta)$ does not depend on k. Hence, the maximal speed of the tail decay of the performance bounds θ is independent of k. Next, we equate the sojourn time bound from Cor. 1 with ε and solve for

$$\tau \le \rho_Q(\theta) + \frac{\ln k + \ln\alpha - \ln\varepsilon}{\theta}, \tag{13}$$

for $\theta > 0$ subject to the stability condition $\rho_Q(\theta) < \rho_A(-\theta)$. Eq. (13) expresses a sojourn time bound that is exceeded at most with probability ε. It exhibits a growth with $\ln k$. The growth is larger for smaller θ corresponding to a higher utilization.

We include an example that demonstrates the quick estimation of the expected value from the sojourn time bound. By integration of the tail of the sojourn time from Cor. 1 we have

$$\begin{aligned}\mathsf{E}[T(n)] &= \int_0^\infty \mathsf{P}[T(n)>\tau]\mathrm{d}\tau \\ &\le \int_0^{\tau^*} \mathrm{d}\tau + k\alpha e^{\theta\rho_Q(\theta)}\int_{\tau^*}^\infty e^{-\theta\tau}\mathrm{d}\tau,\end{aligned}$$

where we used that $\mathsf{P}[T(n) > \tau] \le 1$ to determine $\tau^* = \ln\big(k\alpha e^{\theta\rho_Q(\theta)}\big)/\theta$. It follows that the expected sojourn time

$$\mathsf{E}[T(n)] \le \rho_Q(\theta) + \frac{\ln k + \ln\alpha + 1}{\theta} \tag{14}$$

is also limited by $\ln k$. The result applies for general arrival and service processes and generalizes the finding of $\ln k$ that is obtained in [4], [15] for iid service times.

M|M tasks: In Fig. 3, we consider a fork-join system with $k \ge 1$ parallel servers. Jobs have iid exponential inter-arrival times and are composed of k tasks with iid exponential service times each. The parameters $\rho_A(-\theta)$ and $\rho_Q(\theta)$ for $i \in [1,k]$ are as specified by (7) and (8) where we let $\mu = 1$. We show bounds of the expected sojourn time and sojourn time quantiles τ, where $\mathsf{P}[T(n) > \tau] \le \varepsilon$ and $\varepsilon = 10^{-6}$. The curves show the characteristic logarithmic growth with k. This is also confirmed in simulation results that agree well with the sojourn time bounds. The expectation of the sojourn time (14) is only a rough estimate, as anticipated.

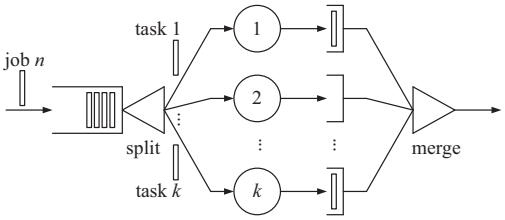

Fig. 4. Split-merge system. Compared to the fork-join system, tasks have an additional synchronization constraint, i.e., the execution of the tasks of a job has to start at the same time.

D. Split-Merge Systems

Split-merge systems, see Fig. 4, are a variant of fork-join systems where all tasks of a job have to start execution simultaneously. If a server i finishes task i of job n, it idles until all tasks $j \in [1,k]$ of that job are finished before any of the tasks of job $n+1$ if any starts.

Lemma 3 (Split-merge system)**.** *Consider a split-merge system with k parallel servers as in Lem. 1. Let $Q_i(n)$ denote the service time of task i of job n where $i \in [1,k]$ and $n \ge 1$. Define for $n \ge m \ge 1$*

$$S(m,n) = \sum_{\nu=m}^{n} \max_{i\in[1,k]}\{Q_i(\nu)\}.$$

The system is an exact $S(m,n)$ server.

Proof: Since all tasks $i \in [1,k]$ of job n start service simultaneously at $V(n)$ and only after all tasks of job $n-1$ have finished service, we have for $n \ge 2$ that

$$V(n) = \max\Big\{A(n), V(n-1) + \max_{i\in[1,k]}\{Q_i(n-1)\}\Big\}, \tag{15}$$

Recursive insertion of (15) with $V(1) = A(1)$ as in Lem. 1, and letting $D(n) = V(n) + \max_{i\in[1,k]}\{Q_i(n)\}$ completes the proof. ∎

Since Lem. 3 proves that the split-merge system satisfies the definition of server Def. 1, the performance bounds of Th. 1 apply, using the parameters of the service process $S(m,n)$ in Lem. 3. In the case of iid service times, the service parameters of Lem. 3 are derived from (6) as

$$\rho_S(\theta) = \frac{1}{\theta}\ln\mathsf{E}\Big[e^{\theta\max_{i\in[1,k]}\{Q_i(1)\}}\Big]. \tag{16}$$

The general problem of split-merge systems is, however, that $\max_{i\in[1,k]}\{Q_i(\nu)\}$ is stochastically increasing with k, with few exceptions such as in the case of identical task service times. The increase implies longer idle times that result in a reduced stability region. As a quick estimate of (16)

$$\rho_S(\theta) \le \frac{1}{\theta}\ln\Big(k\mathsf{E}\Big[e^{\theta Q_1(1)}\Big]\Big)$$

shows that ρ_S has at most a logarithmic growth with the number of parallel servers, resulting in a corresponding reduction of the stability region. A decrease of the stability region with $\ln k$ is also shown in [15], where the authors advise against split-merge implementations based on an in-depth comparison with fork-join systems.

E. Replication Systems

Given k parallel servers, one option to deal with stragglers is the redundant execution of k replicated jobs. In this case, the service time of a job is determined as the minimum of the service times of all its replicas. Once a replica has finished service, all other replicas of the job may or may not be purged. We consider non-purging (k, l) fork-join systems in [1], which includes pure replication systems as a special case if $l = 1$. Systems with replication and purging are elaborated on in [22]. The following lemma says that replication systems with purging satisfy Def. 1. This basic property implies that the performance bounds of Th. 1 hold.

Lemma 4 (Replication system)**.** *Consider a purging replication system with k parallel servers as in Lem. 1. Let $L_i(n)$ denote the service time of replica i of job n where $i \in [1, k]$ and $n \geq 1$. For $n \geq m \geq 1$ define*

$$S(m, n) = \sum_{\nu=m}^{n} \min_{i\in[1,k]} \{L_i(\nu)\}.$$

The system is an exact $S(m, n)$ server.

Proof: First, we note that all replicas of job $n \geq 1$ start service at the same time $V(n)$. This is an immediate consequence of purging all replicas of a job once one replica finishes service. It follows for $n \geq 2$ that

$$V(n) = \max\left\{ A(n), V(n-1) + \min_{i\in[1,k]} \{L_i(n-1)\} \right\}, \quad (17)$$

Recursive insertion of (17) with $V(1) = A(1)$ as in Lem. 1, and letting $D(n) = V(n) + \min_{i\in[1,k]}\{L_i(n)\}$ completes the proof. ∎

Performance bounds follow by insertion of the service parameters of $S(m, n)$ as defined by Lem. 4 into Th. 1. As a detailed evaluation of replication systems with purging and correlated replicas is provided by [22], we only evaluate the simple case of iid replicas with exponential service times L_i with parameter μ. It follows that $\min_{i\in[1,k]}\{L_i(\nu)\}$ is exponential with parameter $k\mu$ so that $\rho_S(\theta)$ follows by substitution of $k\mu$ for μ in (8).

III. Multi-Stage Fork-Join Networks

We contribute a new bound on the growth of end-to-end sojourn times for multi-stage fork-join networks, where we consider h fork-join stages[2] in tandem, each with k parallel servers. We use subscript $i \in [1, k]$ to distinguish the servers of a stage and superscript $j \in [1, h]$ to denote the stages. Note that jobs depart from each fork-join stage in the order of their arrival. The following lemma reproduces a fundamental result of the network calculus [24].

Lemma 5 (Multi-stage fork-join network)**.** *Consider a multi-stage network of h fork-join systems as in Lem. 2 in tandem. Define for $n \geq m \geq 1$*

$$S^{net}(m, n) = \max_{\nu^j : m \leq \nu^1 \leq \nu^2 \leq \ldots \nu^{h-1} \leq n} \{S^1(m, \nu^1) + S^2(\nu^1, \nu^2) + \cdots + S^h(\nu^{h-1}, n)\}.$$

[2]While we focus on multi-stage fork-join networks only, we note that the same analysis applies to networks of split-merge or replication systems.

The fork-join network is an exact $S^{net}(m, n)$ server.

Proof: In a tandem of fork-join systems, the departures of stage j are the arrivals of stage $j + 1$, i.e., $A^{j+1}(n) = D^j(n)$ for $j \in [1, h-1]$. Further, since jobs depart from a fork-join system in the order of their arrival, we have for all $j \in [1, h]$ that $A^j(n) \geq A^j(m) \geq 0$ for $n \geq m \geq 1$. Next, we use that each fork-join stage is an exact server as in Def. 1. We start with $D^h(n) = \max_{\nu^{h-1}\in[1,n]}\{D^{h-1}(\nu^{h-1}) + S^h(\nu^{h-1}, n)\}$ and recursively insert Def. 1 for $j \in [1, h-1]$, to obtain

$$D^h(n) = \max_{m\in[1,n]} \Big\{ \max_{\nu^j : m \leq \nu^1 \leq \cdots \leq \nu^{h-1} \leq n} \big\{A^1(m) + S^1(m, \nu^1) + S^2(\nu^1, \nu^2) + \cdots + S^h(\nu^{h-1}, n)\big\}\Big\},$$

for $n \geq 1$. This proves that $S^{\text{net}}(m, n)$ is an exact server. ∎

Theorem 2 (Multi-stage fork-join network)**.** *Consider a multi-stage fork-join network as in Lem. 5 with arrival and service parameters $(\sigma_A(-\theta), \rho_A(-\theta))$ and $(\sigma_Q(\theta), \rho_Q(\theta))$ as specified by Def. 2. Let the service times at each of the stages be independent. For $n \geq 1$, the end-to-end sojourn time satisfies*

$$\mathsf{P}[T(n) > \tau] \leq k^h \alpha e^{\theta h \rho_Q(\theta)} e^{-\theta\tau},$$

where $\theta > 0$ has to satisfy $\rho_Q(\theta) < \rho_A(-\theta)$ and

$$\alpha = \frac{e^{\theta(\sigma_A(-\theta) + h\sigma_Q(\theta))}}{\left(1 - e^{-\theta(\rho_A(-\theta) - \rho_Q(\theta))}\right)^h}.$$

The proof can be found in the appendix. It uses an established method of the stochastic network calculus [33]; see also the tutorial [37] or the textbook [41].

To see the growth of τ with h and k, we equate the sojourn time bound in Th. 2 with ε and solve for

$$\tau = \sigma_A(-\theta) + h(\sigma_Q(\theta) + \rho_Q(\theta)) + \frac{1}{\theta}\left(h \ln k - h \ln\left(1 - e^{-\theta(\rho_A(-\theta) - \rho_Q(\theta))}\right) - \ln \varepsilon\right)$$

that grows in $\mathcal{O}(h \ln k)$. The result compares to a growth in $\mathcal{O}(h \ln(hk))$ obtained previously in [1]. The improvement is achieved by taking advantage of the statistical independence of the stages, whereas [1] does not make this assumption.

M|M tasks: Fig. 5 shows sojourn time bounds for a multi-stage fork-join network with up to $h = 64$ stages each with $k \in \{1, 2, 4, 8, 16\}$ parallel servers. The inter-arrival and service times are exponential with parameters $\lambda = 0.5$ and $\mu = 1$, respectively, and $\varepsilon = 10^{-6}$ as in Fig. 3. The analytical results in Fig. 5(a) exhibit the same characteristic trends as the simulation results in Fig. 5(b), albeit with less precision due to the inequalities that are involved for each stage. The end-to-end sojourn times show a linear growth with h and a logarithmic growth with k, observable by the equidistantly spaced lines for $k \in \{1, 2, 4, 8, 16\}$.

The simulation results that we depict in Fig. 5(b) are obtained for tasks that have an independent service time at each stage. We also conducted simulations of a multi-stage fork-join network where the service times of tasks are identical at each stage, i.e., not independent. While we omit showing the results, it is interesting to note that the end-to-end sojourn times observed in these simulations grow faster than linearly with h as predicted by $\mathcal{O}(h \ln(hk))$ in [1].

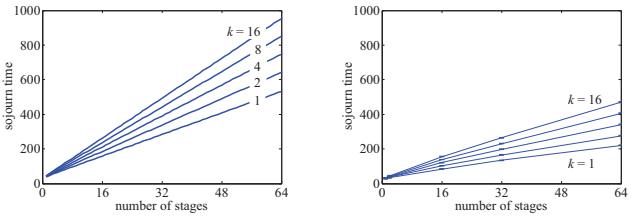

(a) analysis (b) simulation

Fig. 5. Multi-stage fork-join network. Sojourn time bounds grow in $\mathcal{O}(h \ln k)$ for h stages each with k parallel servers.

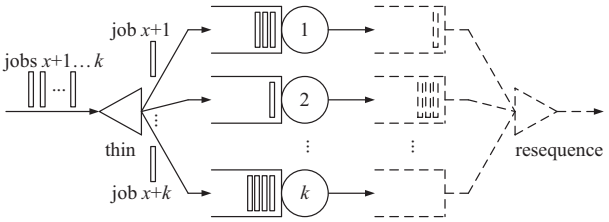

Fig. 6. Multi-server system with thinning. Compared to a fork-join system, jobs are not divided into tasks. Instead, entire jobs are assigned to the parallel servers, e.g., deterministically in round robin order, resulting in a thinning of the arrival process. The multi-server system does not maintain the order of jobs unless an additional resequencing step is added (dashed lines).

IV. Multi-Server Systems with Thinning

In this section, we compare the performance of fork-join systems to that of traditional multi-server systems. An example of a traditional multi-server system with k servers is depicted in Fig. 6. The difference to fork-join systems is that jobs are not divided into tasks that are served in parallel but instead each job is assigned in its entirety to one of the servers. As a consequence, the external arrival process $A(n)$ is divided into k "*thinned*" processes $A_i(m)$. The departures $D_i(m)$ of each server $i \in [1,k]$ may optionally be resequenced in the original order of $A(n)$ to form the departure process $D(n)$.

A. Thinning

First we introduce some notation. While the external arrival process $A(n)$ specifies the arrival time of job n before thinning, the processes $A_i(m)$ denote the arrival time of the mth job of the thinned process at server $i \in [1,k]$. The corresponding service time is $L_i(m)$.

In the case of random thinning, each job is assigned to one of the k servers according to iid discrete (not necessarily uniform) random variables with support $[1,k]$. From the iid property, the mapping of each job to a certain server $i \in [1,k]$ is an independent Bernoulli trial with parameter p_i where $\sum_{i=1}^{k} p_i = 1$. Let $X_i(m)$ denote the number of the job that becomes the mth job that is assigned to server i. It follows that $X_i(m)$ is a sum of m iid geometric random variables with parameter p_i; i.e., $X_i(m)$ is negative binomial. The arrival process at server i is

$$A_i(m) = A(X_i(m)), \tag{18}$$

for $m \geq 1$. Conversely, given jobs $1, 2, \ldots, n$ of the external arrival process, let $Y_i(n)$ denote the number of jobs assigned to server i. It follows that $Y_i(n)$ is binomial with parameter p_i. Further, it holds that $X_i(Y_i(n)) \leq n$ for $n \geq 1$.

In the case of deterministic thinning a round robin assignment of the jobs of an arrival process $A(n)$ to k servers results in the processes $A_i(m)$ as in (18) where

$$X_i(m) = k(m-1) + i, \tag{19}$$

for $m \geq 1$ and $i \in [1,k]$. Given jobs $1, 2, \ldots, n$, the number of jobs that are assigned to server i is

$$Y_i(n) = \left\lceil \frac{n-i+1}{k} \right\rceil, \tag{20}$$

for $n \geq 1$ and $i \in [1,k]$. To see this, note that job n of the external arrival process becomes the $m = \lceil n/k \rceil$th job of server $j = (n-1) \mod k + 1$. Hence, $Y_i(n) = m$ for $i \leq j$ and $Y_i(n) = m-1$ for $i > j$. The same can be verified for (20).

Corollary 2 (Thinning)**.** *Given arrivals with iid inter-arrival times and parameter $\rho_A(-\theta)$ for $\theta > 0$ as in* (5)*. In the case of random thinning with probabilities p_i for $i \in [1,k]$, the thinned arrival processes have parameter*

$$\rho_{A_i}(-\theta) = -\frac{1}{\theta} \ln\left(\frac{p_i e^{-\theta \rho_A(-\theta)}}{1-(1-p_i)e^{-\theta \rho_A(-\theta)}} \right),$$

where $\theta > 0$ so that $e^{-\theta \rho_A(-\theta)} < 1/(1-p_i)$.

In the case of deterministic thinning, for $\theta > 0$ the thinned arrival processes have parameter

$$\rho_{A_i}(-\theta) = k \rho_A(-\theta).$$

Proof: The thinned arrivals are expressed by (18) as a doubly random process that has increments $A_i(\nu, \nu+1) = A(X_i(\nu), X_i(\nu+1))$ for $\nu \geq 1$. Considering iid inter-arrival times, the MGF of the thinned process is

$$\mathsf{M}_{A_i(\nu,\nu+1)}(-\theta) = \mathsf{E}\left[(\mathsf{M}_{A(1,2)}(-\theta))^{X_i(1)} \right], \tag{21}$$

for $\nu \geq 1$. After some reordering, it follows that

$$\mathsf{M}_{A_i(\nu,\nu+1)}(-\theta) = \mathsf{M}_{X_i(1)}(\ln \mathsf{M}_{A(1,2)}(-\theta)). \tag{22}$$

Since $X_i(1)$ is a geometric random variable with MGF $\mathsf{M}_{X_i(1)}(\theta) = p_i e^{\theta}/(1-(1-p_i)e^{\theta})$ for $\theta < -\ln(1-p_i)$, we obtain by insertion of (22) into (5) that

$$\rho_{A_i}(-\theta) = -\frac{1}{\theta} \ln\left(\frac{p_i \mathsf{M}_{A(1,2)}(-\theta)}{1-(1-p_i)\mathsf{M}_{A(1,2)}(-\theta)} \right), \tag{23}$$

where $\theta > 0$ so that $\mathsf{M}_{A(1,2)}(-\theta) < 1/(1-p_i)$.

In the case of deterministic thinning, (21) simplifies to $\mathsf{M}_{A_i(\nu,\nu+1)}(-\theta) = (\mathsf{M}_{A(1,2)}(-\theta))^k$ so that (5) evaluates to $\rho_{A_i}(-\theta) = k\rho_A(-\theta)$. ∎

Since each of the servers of the multi-server system serves its thinned arrivals independently, statistical performance bounds follow straightforwardly by insertion of the arrival parameters from Cor. 2 into Th. 1. The main effect of thinning is captured in the stability condition of Th. 1, which becomes $\rho_L(\theta) < \rho_{A_i}(-\theta)$, where $\rho_{A_i}(-\theta)$ increases with k. Given fixed $\rho_L(\theta)$, the increase of $\rho_{A_i}(-\theta)$ permits larger θ that result in a faster tail decay. Further, the modularity of

this approach implies that the servers may themselves be fork-join systems that can be analyzed by insertion of Cor. 2 into Cor. 1. The results enable a comparison of the performance of fork-join systems with multi-server systems.

First, we investigate how performance bounds for the multi-server system with deterministic thinning grow with k. To achieve comparability with the fork-join results, the multi-server system serves jobs that are composed of k tasks. Jobs are, however, served in their entirety by one of the servers. Given iid task service times with parameter $\rho_Q(\theta)$ for $i \in [1,k]$ as defined by (6), the job service times have parameter $\rho_L(\theta) = k\rho_Q(\theta)$. By insertion of $\rho_L(\theta)$ and $\rho_{A_i}(-\theta)$ from Cor. 2 into Th. 1, we obtain the stability condition $k\rho_Q(\theta) < k\rho_A(-\theta)$. Hence, the maximal θ that achieves the stability condition is independent of k. As a consequence, the waiting time bound from Th. 1 and the speed of the tail decay of the sojourn time bound do not depend on k. Regarding the sojourn time, we equate the bound from Th. 1 with ε and solve for

$$\tau \leq k\rho_Q(\theta) + \frac{\ln\alpha - \ln\varepsilon}{\theta}, \tag{24}$$

for $\theta > 0$ under the stability condition $\rho_Q(\theta) < \rho_A(-\theta)$. Eq. (24) shows a linear growth with k. The result compares to the logarithmic growth established by (13) for the fork-join system. If k is large, the sojourn time of a job at the multi-server system is dominated by its service time that depends linearly on k. The fork-join system avoids this effect, as the tasks of the jobs are served by k servers in parallel.

We also evaluate a hybrid system where the arrivals are divided into a thinned processes that are served by a fork-join sub-systems each with b servers. This type of system was also studied as partial mapping in [15]. The overall system comprises $k = ab$ servers. Given jobs that consist of k tasks, each server of the selected fork-join system has to serve $k/b = a$ tasks per job. Following the same steps as above, we obtain by insertion of Cor. 2 into Cor. 1:

$$\tau \leq a\rho_Q(\theta) + \frac{\ln b + \ln\alpha - \ln\varepsilon}{\theta}, \tag{25}$$

for $\theta > 0$ under the stability condition $a\rho_Q(\theta) < a\rho_A(-\theta)$. As suggested by (25), configurations where $a \sim \ln k$ achieve a logarithmic scaling with k.

$M|E_k$ jobs: For numerical evaluation, we consider arrivals with iid exponential inter-arrival times and jobs with iid Erlang-k service times with parameters λ and μ, respectively. Each job consists of k tasks with iid exponential service times with parameter μ. We consider three system configurations: i) a multi-server system with thinning, ii) a fork-join system, and iii) a hybrid system, all with k servers.

For the multi-server system with thinning (i), performance bounds are derived from Th. 1 using the parameters of the thinned arrival processes as in Cor. 2. Deterministic thinning results in processes $A_i(m)$ where the inter-arrival times are a sum of k exponential random variables; that is, Erlang-k distributed. It follows by insertion of $\rho_A(-\theta)$ from (7) into Cor. 2 that

$$\rho_{A_i}(-\theta) = -\frac{k}{\theta}\ln\left(\frac{\lambda}{\lambda+\theta}\right), \tag{26}$$

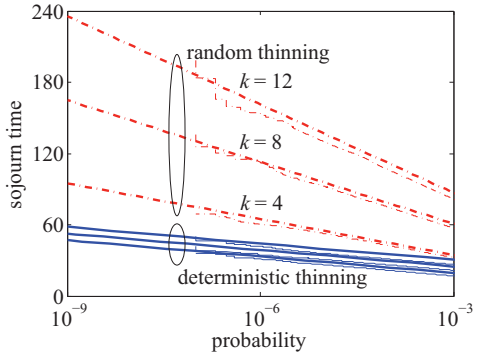

Fig. 7. Multi-server system with thinning. Analytical bounds (thick lines) and simulation results (thin lines down to $\varepsilon = 10^{-7}$). Deterministic thinning outperforms random thinning.

for $\theta > 0$. In the case of random thinning we have by insertion of $\rho_A(-\theta)$ from (7) into Cor. 2 with $p_i = 1/k$ that

$$\rho_{A_i}(-\theta) = -\frac{1}{\theta}\ln\left(\frac{\lambda}{\lambda+k\theta}\right), \tag{27}$$

for $\theta > 0$. In this case, the inter-arrival times of the thinned processes are exponentially distributed with parameter λ/k. Lastly, the Erlang-k service times of the jobs have parameter

$$\rho_L(\theta) = \frac{k}{\theta}\ln\left(\frac{\mu}{\mu-\theta}\right), \tag{28}$$

for $\theta \in (0,\mu)$. For deterministic thinning, the maximal θ that satisfies the stability condition $\rho_L(\theta) \leq \rho_{A_i}(-\theta)$ is $\theta = \mu - \lambda$.

Fig. 7 contrasts the tail decay of deterministic and random thinning for $\mu = 1$, $\lambda = 0.5$, and $k \in \{4, 8, 12\}$. The speed of the tail decay does not depend on k for deterministic thinning. In contrast, for random thinning the tail decay becomes slower with increasing k. Deterministic thinning generally outperforms random thinning. In the following comparison we include only deterministic thinning.

The fork-join system (ii) is the same as already evaluated in Fig. 3. For the hybrid system (iii), the thinned arrival process has Erlang-a inter-arrival times and since each job has k tasks that are divided among b servers, each server has to serve $k/b = a$ exponential tasks, resulting in sum in an Erlang-a service time. We choose $a = \log_2 k$ so that $b = k/\log_2 k$. Performance bounds for this system are derived from Cor. 1 by insertion of the parameters from Cor. 2.

Fig. 8(a) and Fig. 8(b) compare the performance of the three systems for $\mu = 1$, $\lambda \in \{0.5, 0.7, 0.8\}$, and $\varepsilon = 10^{-6}$. Clearly, the sojourn time bounds of the multi-server system with deterministic thinning grow linearly with k, as established by (24). For large k the sojourn time of a job is dominated by its service time, so that the curves that are depicted for different arrival rates λ converge eventually. The fork-join system mitigates the impact of large jobs by serving the tasks in parallel. It achieves a scaling in $\ln k$, see (13), that is due to the synchronization constraint of the join operation. The fork-join system mostly outperforms the multi-server system, except in the case of large λ and small k. The reason is that in a fork-join system the occurrence of a large task blocks all subsequent tasks of that server so the respective jobs cannot

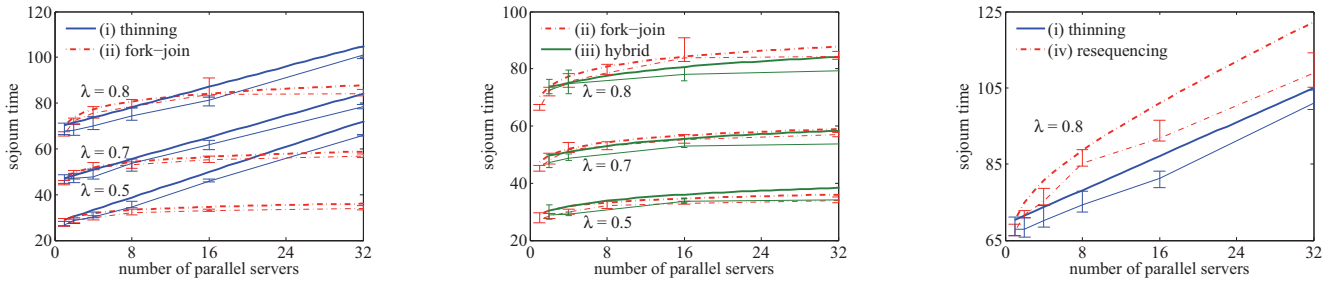

(a) thinning vs. fork-join (b) fork-join vs. hybrid (c) thinning with resequencing

Fig. 8. Comparison of different multi-server configurations. Analytical bounds (thick lines) and simulation results (thin lines). Sojourn time bounds grow linearly with k in the case of multi-server systems with thinning (i) and with $\ln k$ for fork-join systems (ii), each with k servers. The hybrid system (iii) is optimized to achieve a growth with $\ln k$, as well. Resequencing (iv) adds another delay that grows with $\ln k$ to the multi-server system with thinning (i).

complete the join operation. In contrast, in a multi-server system there is no synchronization constraint, so subsequent jobs that are served by other servers can finish service earlier. A similar argument applies for the hybrid system that achieves a logarithmic scaling with k, see (25), and outperforms the fork-join system if λ is large.

B. Resequencing

Unlike fork-join systems, multi-server systems do not guarantee that jobs depart in their order of arrival. An optional resequencing step is depicted in Fig. 6. As the resequencing step does not affect the waiting time of a job, we state the following corollary only for the sojourn time.

Corollary 3 (Thinning and resequencing)**.** *Consider a multi-server system with k parallel servers as in Lem. 1, thinning, and resequencing. The thinned arrivals have parameter $\rho_{A_i}(-\theta)$ as specified by Cor. 2 and the jobs have service parameters $(\sigma_L(\theta), \rho_L(\theta))$ as specified by Def. 2. For $n \geq 1$, the sojourn time satisfies*

$$\mathsf{P}[T(n) > \tau] \leq k\alpha e^{\theta\rho_L(\theta)} e^{-\theta\tau}.$$

In the case of GI|G arrival and service processes, the free parameter $\theta > 0$ has to satisfy $\rho_L(\theta) < \rho_{A_i}(-\theta)$ and

$$\alpha = \frac{e^{\theta\sigma_L(\theta)}}{1 - e^{-\theta(\rho_{A_i}(-\theta) - \rho_L(\theta))}}.$$

In the case of GI|GI arrival and service processes, $\theta > 0$ has to satisfy $\rho_L(\theta) \leq \rho_{A_i}(-\theta)$ and $\alpha = 1$.

Proof: Given the departure processes of each of the servers $D_i(n)$ for $i \in [1,k]$, the combined in-sequence departure process for $n \geq 0$ is

$$D(n) = \max_{i\in[1,k]} \{D_i(Y_i(n))\}, \tag{29}$$

where $D_i(0) = 0$ by convention. Note that in evaluating (29) one only has to verify the departure of job $Y_i(n)$ from each server $i \in [1,k]$, since the departure of job $Y_i(n)$ from server i implies the departure of all jobs $\nu \in [1, Y_i(n)]$ of the same server. I.e., $D_i(\nu) \leq D_i(Y_i(n))$ for $\nu \in [1, Y_i(n)]$, since each server implements first-in first-out order.

Next, we use Lem. 1 for each server $i \in [1,k]$ to obtain $D_i(Y_i(n)) = \max_{m\in[1,Y_i(n)]}\{A_i(m) + S_i(m, Y_i(n))\}$. By insertion into (29) we have

$$D(n) = \max_{i\in[1,k]} \left\{ \sup_{m\in[1,Y_i(n)]} \{A_i(m) + S_i(m, Y_i(n))\} \right\}, \tag{30}$$

for $n \geq 1$, where $\sup\{\emptyset\} = 0$. We estimate the sojourn time $T(n) = D(n) - A(n) \leq D(n) - A_i(Y_i(n))$ for $n \geq 1$, where we used that $A(n) \geq A_i(Y_i(n))$ for $i \in [1,k]$. By insertion of (30), it follows for $n \geq 1$ that

$$T(n) \leq \max_{i\in[1,k]} \left\{ \sup_{m\in[1,Y_i(n)]} \{S_i(m, Y_i(n)) - A_i(m, Y_i(n))\} \right\}.$$

By the union bound $\mathsf{P}[T(n) > \tau] \leq \sum_{i=1}^{k} \mathsf{P}[T_i(n) > \tau]$, where $T_i(n) \leq \sup_{m\in[1,Y_i(n)]}\{S_i(m, Y_i(n)) - A_i(m, Y_i(n))\}$. Finally, we estimate $\mathsf{P}[T_i(n) > \tau]$ using Th. 1. ■

To evaluate the effect of k, we consider the case where the service time grows with k, expressed as $\rho_L(\theta) = k\rho_Q(\theta)$, as before. We investigate deterministic thinning and equate the sojourn time bound from Cor. 3 with ε to solve for

$$\tau \leq k\rho_Q(\theta) + \frac{\ln k + \ln\alpha - \ln\varepsilon}{\theta} \tag{31}$$

for $\theta > 0$ under the stability condition $k\rho_Q(\theta) < k\rho_A(-\theta)$. Eq. (31) shows two effects: a linear growth with k that is due to the increase of the job service time, as also observed for the multi-server system in (24), and a logarithmic term that is due to resequencing.

M|E$_k$ jobs: Fig. 8(c) shows a numerical comparison of multi-server systems with deterministic thinning, with resequencing (iv) and without resequencing (i). We use the same parameters as above. The results clearly show the additional logarithmic delay due to resequencing. Compared to the fork-join system, resequencing consumes the advantage that multi-server systems with thinning showed for large λ in Fig. 8(a).

V. Non-Idling Systems

A major drawback of the basic fork-join system, as well as the multi-server system with thinning, is that servers may idle while tasks (or jobs, respectively) are queued at other servers. This is due to the early and static assignment of

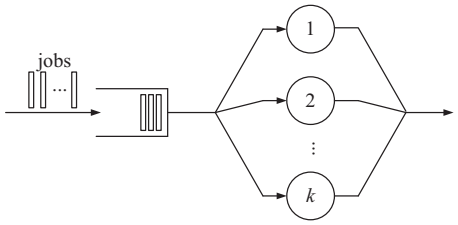

Fig. 9. Single-queue multi-server system. The system is non-idling: once a server finishes a job, the head of the line job is assigned to that server.

tasks to servers at their time of arrival, that necessitates an individual queue at each server. One advantage of the early assignment is that data can be prefetched at the respective server during the waiting time. If a late assignment of tasks to servers is possible, an implementation with a single queue avoids idling; that is, whenever a server becomes idle, the system assigns the task (or job) at the head of the queue to that server. Compared to multi-queue systems, where each server has an individual queue, single-queue systems use additional feedback information from the servers to perform the late assignment of tasks. We assume that feedback and task assignment are instantaneous. Otherwise, e.g., in a distributed implementation, each task incurs an overhead time in addition to its service time. A further difficulty of single-queue systems is that jobs can depart in an order that differs from the order of their arrival, i.e., $D(n) \not\geqslant D(n-1)$. This implies also that the waiting time of job n cannot simply be determined from $D(n-1)$ as in (4) and Th. 1, nor from $D(n-k)$. In the following subsections, we first provide performance bounds for single-queue multi-server systems, i.e., G|M|k queues. The bounding method that we derive here, extends to single-queue fork-join systems that we present next.

A. Single-Queue Multi-Server Systems

We consider single-queue multi-server systems without a resequencing constraint. An example is shown in Fig. 9. First, we prove that the system satisfies the definition of max-plus server and based thereon we derive performance bounds.

Lemma 6 (Single-queue multi-server system)**.** *Consider a single-queue multi-server system with k parallel servers as in Lem. 1. Let $L(n)$ denote the service time of job n for $n \geq 1$. Given that all servers are busy after job n starts service at $V(n)$, define $Z(n)$ to be the time until the next server becomes idle. Otherwise let $Z(n) = 0$. Define for $n \geq m \geq 1$*

$$S(m,n) = L(n) + \sum_{\nu=m}^{n-1} Z(\nu).$$

i) The system is an exact $S(m,n)$ server.

ii) Given that the jobs have iid exponential service times with parameter μ. The non-zero elements of $Z(n)$ are iid exponential random variables with parameter $k\mu$.

iii) Replace the zero elements of $Z(n)$ by iid exponential random variables with parameter $k\mu$ and compute $S(m,n)$ as above. The system is an $S(m,n)$ server.

Proof: Using the definition of $Z(n)$, it holds for $n \geq 2$ that

$$V(n) = \max\{A(n), V(n-1) + Z(n-1)\}. \tag{32}$$

Further, $V(1) = A(1)$ and since $Z(n) = 0$ for $n \in [1, k-1]$ we have $V(n) = A(n)$ also for $n \in [2, k]$. By recursive insertion of (32) we obtain for $n \geq 1$ that

$$V(n) = \max_{m\in[1,n]} \left\{ A(m) + \sum_{\nu=m}^{n-1} Z(\nu) \right\}. \tag{33}$$

Since $D(n) = V(n) + L(n)$ we have with (33) that $D(n) = \max_{m\in[1,n]}\{A(m) + S(m,n)\}$, which proves the first part.

Next, we consider iid exponential service times with parameter μ and investigate the distribution of $Z(n)$ for $n \geq 1$. Given all servers are busy after the start of service of job n at $V(n)$. This implies that there are another $k-1$ jobs with indices smaller n that are already in service at $V(n)$. Due to the memorylessness of the exponential distribution, the residual service time of each of these jobs as well as the service time of job n are iid exponential random variables with parameter μ. Since the minimum of k iid exponential random variables with parameter μ is an exponential random variable with parameter $k\mu$, it follows that the time until the next server becomes idle is exponentially distributed with parameter $k\mu$.

For the last part, we use that exponential random variables are non-negative. If we replace all $Z(n)$ that are zero by iid exponential random variables with parameter $k\mu$, (33) becomes $V(n) \leq \max_{m\in[1,n]}\{A(m) + \sum_{\nu=m}^{n-1} Z(\nu)\}$ for $n \geq 1$, and consequently $D(n) \leq \max_{m\in[1,n]}\{A(m) + S(m,n)\}$. ∎

Considering iid exponential service times $L(n)$ with parameter μ, we have $\rho_L(\theta)$ for $\theta < \mu$ as in (8). With Lem. 6 (iii), $Z(n)$ is composed of iid exponential random variables with parameter $k\mu$. Invoking (8) with parameter $k\mu$ gives

$$\rho_Z(\theta) = \frac{1}{\theta} \ln\left(\frac{k\mu}{k\mu - \theta}\right), \tag{34}$$

for $\theta < k\mu$. With (8) and (34) the MGF of $S(m,n)$ in Lem. 6 (iii) is $\mathsf{E}\left[e^{\theta S(m,n)}\right] = e^{\theta(\rho_L(\theta)+\rho_Z(\theta)(n-m))}$ for $\theta < \mu$. Hence, $S(m,n)$ satisfies Def. 2 with parameters $\sigma_S(\theta) = \rho_L(\theta) - \rho_Z(\theta)$ and $\rho_S(\theta) = \rho_Z(\theta)$.

As verified by Lem. 6 (iii), the single-queue multi-server system is an $S(m,n)$ server, where $S(m,n)$ is composed of independent increments. Hence, a sojourn time bound can be derived from Th. 1 using the parameters $(\sigma_S(\theta), \rho_S(\theta))$ defined above. However, the waiting time bound of Th. 1 uses a definition of waiting time $W(n) = [D(n-1) - A(n)]^+$ that does not apply to the single-queue multi-server system, where the departure of job $n-1$ does not generally mark the start of the service of job n. The following theorem first formulates the waiting time for single-queue multi-server systems and in a second step derives a sojourn time bound from the waiting time. The derivation avoids the technical limitation $\theta < \mu$ that applies if (8) is inserted into Th. 1 and thus enables tighter bounds.

Theorem 3 (Single-queue multi-server system)**.** *Consider a single-queue multi-server system as in Lem. 6, with arrival parameters $(\sigma_A(-\theta), \rho_A(-\theta))$ as specified by Def. 2, iid*

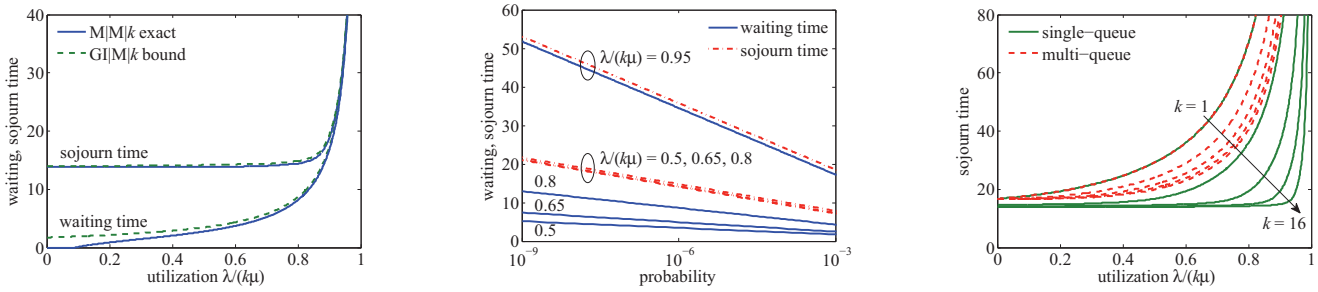

(a) bound vs. exact result (b) tail decay (c) single- vs. multi-queue system

Fig. 10. Single-queue multi-server system. (a) The bounds agree closely with the exact result. (b) For high utilization the sojourn time is dominated by the waiting time, otherwise by the service time. (c) Comparison of single-queue and multi-queue systems with thinning for $k = \{1, 2, 4, 8, 16\}$. The sojourn time decreases with k. The improvement is, however, significantly larger in the case of the single-queue system that can sustain a utilization close to one if k is large.

exponential job service times with parameter μ, and parameter $\rho_Z(\theta)$ given by (34). *For $n \geq 1$, the sojourn time satisfies*

$$\mathsf{P}[T(n) > \tau] \leq \alpha \frac{\mu}{\mu - \theta} \left(e^{-\theta\tau} - e^{-\mu\tau} \right) + e^{-\mu\tau},$$

and the waiting time

$$\mathsf{P}[W(n) > \tau] \leq \alpha e^{-\theta\tau}.$$

In the case of G|M arrival and service processes, the free parameter $0 < \theta < k\mu, \theta \neq \mu$ has to satisfy $\rho_Z(\theta) < \rho_A(-\theta)$ and

$$\alpha = \frac{e^{\theta\sigma_A(-\theta)}}{1 - e^{-\theta(\rho_A(-\theta) - \rho_Z(\theta))}}.$$

In the special case of GI|M arrival and service processes, $0 < \theta < k\mu, \theta \neq \mu$ has to satisfy $\rho_Z(\theta) \leq \rho_A(-\theta)$ and $\alpha = 1$.

Proof: With Lem. 6 (iii) and (33), the waiting time $W(n) = V(n) - A(n)$ for $n \geq 1$ is estimated as

$$W(n) \leq \max_{m \in [1,n]} \left\{ \sum_{\nu=m}^{n-1} Z(\nu) - A(m,n) \right\},$$

where $Z(n)$ for $n \geq 1$ are iid exponential random variables with parameter $k\mu$. The derivation of the statistical waiting time bound closely follows the proof of Th. 1 in the appendix and is omitted.

To derive a sojourn time bound we use that $D(n) = V(n) + L(n)$ so that $T(n) = D(n) - A(n)$ can be expressed as $T(n) = W(n) + L(n)$, where we substituted $W(n) = V(n) - A(n)$. We use the waiting time bound from Th. 3 to estimate the waiting time CDF[3] as $F_W(\tau) = \mathsf{P}[W(n) \leq \tau] \geq 1 - \alpha e^{-\theta\tau}$. By convolution with the exponential job service time PDF $f_L(\tau) = \mu e^{-\mu\tau}$ for $\tau \geq 0$ we obtain the CDF of the sojourn time as

$$F_T(\tau) = \int_0^\tau F_W(\tau - x) f_L(x) dx$$

that evaluates for $\theta \neq \mu$ to

$$F_T(\tau) \geq 1 - e^{-\mu\tau} - \alpha \frac{\mu}{\mu - \theta} (e^{-\theta\tau} - e^{-\mu\tau}),$$

which completes the proof. ∎

[3]We note that for $\alpha > 1$, a tighter bound can be derived, using that $F_W(\tau) = 0$ for $\tau < -\ln(1/\alpha)/\theta$. We omit the details for notational brevity.

M|M|k Queue: For numerical evaluation, we use jobs with exponential inter-arrival times with parameter λ and exponential service times with parameter $\mu = 1$. In this case, the system becomes the well-known M|M|k queue for which there is an exact solution. We note that Th. 3 is not limited to exponential arrivals. Fig. 10(a) compares the waiting time and the sojourn time bounds with the exact results for $k = 8$, $\varepsilon = 10^{-6}$, and different utilizations defined as $\lambda/(k\mu)$. The exact waiting time distribution of the M|M|k queue is $\mathsf{P}[W(n) > \tau] = P_k e^{-(k\mu - \lambda)\tau}$ where P_k is the probability of waiting, i.e., the probability that k or more jobs are in the system [42]. The bounds from Th. 3 are obtained by insertion of $\rho_A(-\theta)$ from (7). From the stability condition $\rho_Z(\theta) \leq \rho_A(-\theta)$ we find $\theta \leq k\mu - \lambda$. By the choice of maximal θ, the waiting time bound $\mathsf{P}[W(n) > \tau] \leq e^{-(k\mu - \lambda)\tau}$ exhibits the same exponential speed of decay and differs by the prefactor P_k that approaches one for high utilization. Fig. 10(a) confirms the good agreement of the waiting time bound and shows visible differences only in the case of low utilization.

Regarding the sojourn time bound in Fig. 10(a), we generally have good agreement. Here, two effects can be distinguished. These are expressed by the two parts of the sojourn time bound in Th. 3 that decay as $e^{-\theta\tau}$ and $e^{-\mu\tau}$, respectively. In the case of high utilization, the sojourn time is dominated by the waiting time that decays with $e^{-\theta\tau}$ where $\theta = k\mu - \lambda$. Otherwise, if the utilization satisfies $\lambda/(k\mu) < 1 - 1/k$ (that is 0.875 for $k = 8$ here), it follows that $\theta > \mu$ so that the waiting time decays quickly and the sojourn time is mostly due to the service time of the job itself that decays slower with $e^{-\mu\tau}$. Fig. 10(b) details the effect, again for $k = 8$. For utilizations below 0.875, the waiting time decays faster than the service time so that the sojourn time changes only marginally if the utilization is increased from 0.5 to 0.8. In contrast, once the utilization exceeds 0.875, the waiting time dominates.

Fig 10(c) evaluates the sojourn time of the single-queue multi-server system for $k = \{1, 2, 4, 8, 16\}$ with respect to the multi-queue system with thinning. For comparability, we use the same technique to derive the sojourn time from the waiting time as in Th. 3 for both systems. While for $k = 1$ the systems are identical, the single-queue multi-server system outperforms the multi-queue system for larger k. Given a target

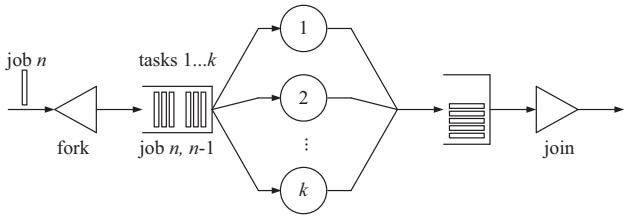

Fig. 11. Single-queue fork-join system. The system is non-idling: once a server finishes a task, the head of the line task is assigned to that server. The join operation uses a random access buffer to complete jobs (possibly out of sequence) once all tasks are finished.

sojourn time bound, the single-queue system can sustain a significantly higher utilization.

G|D|k Queue: While single-queue multi-server systems show a significant advantage if the service times are exponential, we note that this is not generally the case. An example are deterministic service times, i.e., $L(n) = L$ for $n \geq 1$. In this case, the single-queue multi-server system is governed by $V(n) = \max\{A(n), V(n-k) + L\}$ for $n > k$ and $V(n) = A(n)$ for $n \in [1, k]$. This is an immediate consequence of the deterministic service times, which ensure that all jobs finish service in the order of their arrival. Considering the multi-queue multi-server system with deterministic thinning, job n is assigned to server $i = (n-1) \mod k + 1$ following job $n-k$. Hence job n starts service at $V(n) = \max\{A(n), V(n-k) + L\}$ for $n > k$ and $V(n) = A(n)$ for $n \in [1, k]$. This shows that the single-queue und the multi-queue multi-server system perform identical in case of deterministic service times.

B. Single-Queue Fork-Join Systems

In a single-queue fork-join system, jobs are composed of k tasks that are stored in a single-queue. Once any of the k parallel servers becomes idle, it fetches the next task from the head of the queue. An example is shown in Fig. 11. As tasks may finish service out-of-sequence, the join operation uses a buffer with random access to complete a job immediately once all of its tasks are finished, regardless of the order of arrival. This implies that $D(n) \not\geq D(n-1)$. Our analysis of single-queue fork-join systems follows the same essential steps as in the case of the single-queue multi-server systems with the additional resynchronization constraint of the join operation.

Lemma 7 (Single-queue fork-join system)**.** *Consider a single-queue fork-join system with k parallel servers as in Lem. 1. Let $Q_i(n)$ denote the service time of task i of job n for $n \geq 1$ and $i \in [1, k]$. Given that all servers are busy after task i of job n starts service at $V_i(n)$, define $Z_i(n)$ to be the time until the next server becomes idle. Otherwise let $Z_i(n) = 0$. Define for $n \geq m \geq 1$*

$$S(m,n) = \max_{i\in[1,k]} \left\{ Q_i(n) + \sum_{j=1}^{i-1} Z_j(n) + \sum_{\nu=m}^{n-1} \sum_{j=1}^{k} Z_j(\nu) \right\}.$$

i) The system is an exact $S(m,n)$ server.

ii) Given that the tasks have iid exponential service times with parameter μ. The non-zero elements of $Z_i(n)$ are iid exponential random variables with parameter $k\mu$.

iii) Replace the zero elements of $Z_i(n)$ by iid exponential random variables with parameter $k\mu$ and compute $S(m,n)$ as above. The system is an $S(m,n)$ server.

Proof: Using the definition of $Z_i(n)$, it holds for $n \geq 1$ and $i \in [2, k]$ that

$$V_i(n) = \max\{A(n), V_{i-1}(n) + Z_{i-1}(n)\}, \tag{35}$$

and for $n \geq 2$ and $i = 1$

$$V_1(n) = \max\{A(n), V_k(n-1) + Z_k(n-1)\}. \tag{36}$$

Further, $V_1(1) = A(1)$ and since $Z_i(1) = 0$ for $i \in [1, k-1]$ we have $V_i(1) = A(1)$ also for $i \in [2, k]$. By recursive insertion of (35) and (36) we obtain for $n \geq 1$ and $i \in [1, k]$ that

$$V_i(n) = \max_{m\in[1,n]} \left\{ A(m) + \sum_{j=1}^{i-1} Z_j(n) + \sum_{\nu=m}^{n-1} \sum_{j=1}^{k} Z_j(\nu) \right\}. \tag{37}$$

Above we used that $Z_i(n)$ is non-negative to reduce the number of terms that are evaluated by the maximum operator. Given $V_i(n)$, task i of job n finishes service after another $Q_i(n)$ units of time at $D_i(n) = V_i(n) + Q_i(n)$. Finally, job n is completed once all of its tasks have finished service at $D(n) = \max_{i\in[1,k]}\{V_i(n) + Q_i(n)\}$. Inserting (37) and reordering the maxima proves the first part.

The proof of the remaining parts is a notational extension of the proof of Lem. 6 that considers tasks instead of jobs. ■

Considering exponential service times $Q_i(n)$ with parameter μ, we have $\rho_Q(\theta)$ as in (8) for $\theta < \mu$. With Lem. 7 (iii), $Z_i(n)$ is composed of iid exponential random variables with parameter $k\mu$ that are characterized by $\rho_Z(\theta)$ given by (34) for $\theta < k\mu$. The MGF of $S(m,n)$ in Lem. 7 (iii) for $\theta < \mu$ is

$$\begin{aligned} \mathsf{E}\left[e^{\theta S(m,n)}\right] &\leq \sum_{i=1}^{k} e^{\theta(\rho_Q(\theta)+\rho_Z(\theta)((n-m)k+i-1))} \\ &= \beta e^{\theta(\rho_Q(\theta)+k\rho_Z(\theta)(n-m))}, \end{aligned}$$

where

$$\beta = \frac{e^{\theta k \rho_Z(\theta)} - 1}{e^{\theta\rho_Z(\theta)} - 1}. \tag{38}$$

Hence, $S(m,n)$ satisfies Def. 2 with parameters $\sigma_S(\theta) = \rho_Q(\theta) - k\rho_Z(\theta) + \ln\beta/\theta$ and $\rho_S(\theta) = k\rho_Z(\theta)$.

Theorem 4 (Single-queue fork-join system)**.** *Consider a single-queue fork-join system as in Lem. 7, with arrival parameters $(\sigma_A(-\theta), \rho_A(-\theta))$ as specified by Def. 2, iid exponential task service times with parameter μ, parameter $\rho_Z(\theta)$ as specified by* (34)*, and β as given in* (38)*. For $n \geq 1$, the sojourn time satisfies*

$$\mathsf{P}[T(n) > \tau] \leq \alpha\beta \frac{\mu}{\mu - \theta} \left(e^{-\theta\tau} - e^{-\mu\tau}\right) + e^{-\mu\tau},$$

and the waiting time of task i

$$\mathsf{P}[W_i(n) > \tau] \leq \alpha e^{\theta(i-1)\rho_Z(\theta)} e^{-\theta\tau}.$$

In the case of G|M arrival and service processes, the free parameter $0 < \theta < k\mu, \theta \neq \mu$ has to satisfy $k\rho_Z(\theta) < \rho_A(-\theta)$ and

$$\alpha = \frac{e^{\theta\sigma_A(-\theta)}}{1 - e^{-\theta(\rho_A(-\theta) - k\rho_Z(\theta))}}.$$

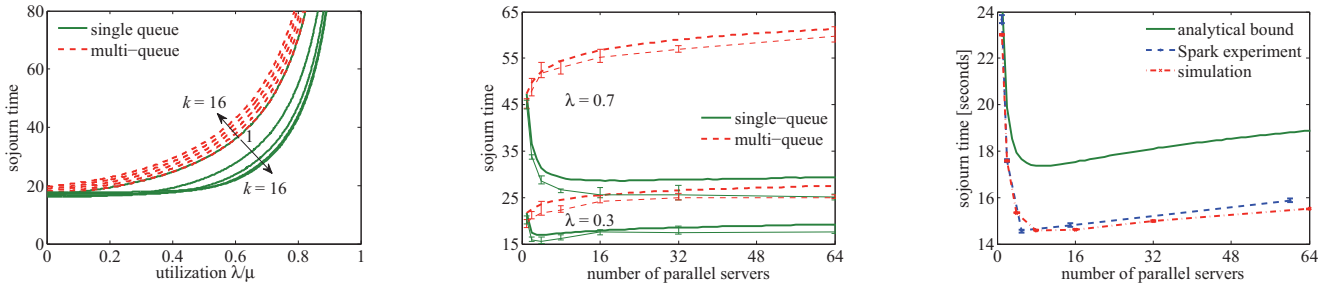

(a) $k = \{1, 2, 4, 8, 16\}$, $\varepsilon = 10^{-6}$ (b) $\lambda = \{0.3, 0.7\}$, $\varepsilon = 10^{-6}$ (c) Spark, $\lambda = 0.7$, $\varepsilon = 10^{-3}$

Fig. 12. (a) and (b) Comparison of single-queue and multi-queue fork-join systems. Analytical bounds (thick lines) and simulation results (thin lines). While the sojourn time of the multi-queue system grows with $\ln k$, the single-queue system achieves a significant improvement due to load-balancing if k is increased, particularly in the case of high utilization. The additional gain diminishes for large k and is eventually consumed by the synchronization constraint of the join operation (parameter β in Th. 4). (c) Spark experiment. The bound predicts the trend of the sojourn time with k of the Spark system.

In the special case of GI|M arrival and service processes, $0 < \theta < k\mu, \theta \neq \mu$ *has to satisfy* $k\rho_Z(\theta) \leq \rho_A(-\theta)$ *and* $\alpha = 1$.

We note that the waiting time of the task of job n that starts service last is simply the waiting time of task k of job n as all other tasks start service before, i.e., no maximum as in the multi-queue fork-join system is needed.

Proof: Similar to the proof of Th. 3, we start with the waiting time that is expressed as $W_i(n) = V_i(n) - A(n)$ for task $i \in [1, k]$ of job $n \geq 1$. With Lem. 7 (iii) and (37) we have for $n \geq 1$ that

$$W_i(n) \leq \max_{m \in [1,n]} \left\{ \sum_{j=1}^{i-1} Z_j(n) + \sum_{\nu=m}^{n-1} \sum_{j=1}^{k} Z_j(\nu) - A(m,n) \right\},$$

where $Z_i(n)$ for $n \geq 1$ and $i \in [1, k]$ are iid exponential random variables with parameter $k\mu$. The derivation of the statistical waiting time bound closely follows the proof of Th. 1 in the appendix and is therefore omitted.

A sojourn time bound for each task $i \in [1, k]$ of job n follows from its waiting time bound by convolution with the exponential task service time PDF as in the proof of Th. 3. Finally, estimating the maximum sojourn time of all tasks $i \in [1, k]$ of job n by use of the union bound leads to parameter β defined by (38). ∎

M|M tasks: We compare the single-queue fork-join system with the multi-queue system from Sec. II-C. Jobs have iid exponential inter-arrival times and are composed of k tasks with iid exponential service times. The parameters $\rho_A(-\theta)$ and $\rho_Q(\theta)$ are as specified by (7) and (8), respectively, where we let $\mu = 1$. For comparability, we use the same technique to derive the sojourn time from the waiting time as in Th. 4 for both systems. In Fig. 12(a), we fix $\varepsilon = 10^{-6}$ and show the impact of the utilization λ/μ for different $k \in \{1, 2, 4, 8, 16\}$. For $k = 1$ the single-queue and the multi-queue system are identical and the sojourn time bounds from Cor. 1 and Th. 4 agree. For increasing k, the sojourn time bound of the multi-queue fork-join system shows logarithmic growth with k; i.e., the lines are equally spaced. This effect is due to the synchronization constraint of the join operation. In contrast, the sojourn time bounds of the single-queue fork-join system improve with k with decreasing gain. Here, two opposing effects are superimposed: 1.) a gain due to load balancing is achieved by the single-queue system if k is increased, most visible for medium to high utilization; 2.) the synchronization constraint of the join operation, similar to the case of the multi-queue fork-join system. Fig. 12(b) depicts the effects for fixed $\lambda = 0.3$ and $\lambda = 0.7$, respectively, $\varepsilon = 10^{-6}$, and varying k. For small λ the gain of the single-queue system that is achieved by load-balancing is small. For intuition, if all servers of the two systems are idle at the time of a job arrival, the single-queue and the multi-queue system perform identically and the sojourn time is determined by the task with the maximal service time. For large λ the gain of load-balancing becomes more significant as k is increased, but for large k the synchronization constraint of the join operation eventually consumes the gain.

The default scheduler of Apache Spark is a prominent implementation of a single-queue system. Fig. 12(c) shows results for $\lambda = 0.7$, $\mu = 1$, and $\varepsilon = 10^{-3}$ from experiments on a live Spark cluster. The units are in seconds. For further details on the experiment see the appendix. Our simulation results match the Spark measurements almost perfectly. Similarly, the sojourn time bound from Th. 4 clearly shows the trend that is observed for Spark if k is increased. First the gain due to load-balancing dominates, and that is later consumed by the synchronization constraint.

VI. Conclusions

We formulated a general model of fork-join systems in max-plus system theory and derived performance bounds for fork-join networks with h independent fork-join stages each with k parallel G|G|1 servers. The bounds were shown to scale in $\mathcal{O}(h \ln k)$ and compare to the previous result $\mathcal{O}(h \ln(hk))$ that does not take advantage of independence. We performed a detailed comparison of essential configurations of multi-server systems. We included an analysis of single-queue multi-server as well as single-queue fork-join systems that are non-idling as opposed to corresponding multi-queue implementations. We found that the single-queue systems achieve a fundamental performance gain that is due to load-balancing and possible overtaking of jobs. Since jobs can depart out of sequence, the multi-stage analysis of single-queue systems

is more difficult and remains as an open research question. We included reference results, mostly from simulation, as well as measurements obtained from Spark experiments which show that the analytical bounds closely predict the actual performance of systems.

Appendix A: Simulation and Experiments

Forkulator: `Forkulator` is an event-driven simulator written in Java[4]. A user can choose from single-queue, multi-queue, multi-stage, and (k, l) fork-join systems, as well as multi-server systems with thinning. The arrival and service processes can also be specified as constant rate, exponential, Erlang, normal, or Weibull, and optionally regulated through a leaky bucket. The simulator samples jobs at a user-configurable interval and records the sojourn, waiting, and service times.

To ensure that our samples are close to iid, in our experiments we sample every 100th job. We chose that interval based on an empirical analysis of the autocorrelation of sojourn times in trial simulations. The confidence intervals plotted on all of the simulation results are 68.2% and are computed for the quantile statistics using the method described in [43, Sec. 2.2.2]. In all cases we ran at least 10^9 iterations, giving about 10^7 samples, which is enough to estimate the $1-10^{-6}$ quantile and its confidence interval.

Spark: Spark [3] is a popular data processing engine that implements the map-reduce model. It is part of the Apache Hadoop ecosystem. Within a Spark program one can execute map and reduce style parallelized operations, and with some APIs one can control the degree of parallelism used. This allows us to effectively create jobs containing a controllable number of tasks, and set these tasks to execute for controllable lengths of time. Therefore we can submit a Spark program that is allocated k cores and executes jobs with k' tasks, and we can draw the execution times of these tasks from any distribution with non-negative support. The default task scheduler in Spark puts the tasks in each job in a first-in first-out queue and distributes them to executors as they become available. This mode of operation corresponds to the non-idling, single-queue fork-join model in Sec. V-B.

We have set up a stand-alone Spark cluster on four 24-core servers. By running each Spark slave in a Docker container and limiting each slave to a single core, we are able to run at least 15 single-core workers on each node and effectively emulate a 60-node cluster. We configure the host so that each container has its own IP address. This is not a perfect emulation of a 60-node cluster, since the Docker containers on each node share a network stack, but the jobs in our experiments produce very little network traffic and have a trivial reduce stage, so for our purposes this inaccuracy is inconsequential.

Our `spark-arrivals` program[5] produces jobs with inter-arrival times drawn from an exponential distribution. For each job it spawns a new thread which submits a Spark job, parallelized with k tasks ("slices" in Spark terminology). Running each job in a separate thread is necessary because `parallelize()` is blocking. Without multi-threading, each job would only start after the previous job departed, making the system more like a single-queue split-merge system.

The execution time of each task is drawn from either an exponential or Erlang-k distribution. The tasks generate random numbers and check how long they have been executing. Therefore the time the tasks spend in their execution loops has the desired distribution. However, there is some overhead associated with executing the tasks, changing the service time distribution somewhat. The main components of this are *task deserialization time* and *scheduler delay*. Task deserialization time is the time needed to distribute the task and associated data to the executor. *Scheduler delay* is not as well documented, but appears to be how Spark accounts for any other overhead that is not execution time or deserialization time. These two components combined tended to be about 6ms. In all our experiments the tasks had a mean service time of 1 second, so this overhead was negligible. Since our reduce stage was trivial, the shuffle and result serialization components of the task overhead were always effectively zero. For our Spark experiments we ran $5 \cdot 10^8$ iterations and report the $1 - 10^{-3}$ quantile.

[4]Software available at https://github.com/brentondwalker/forkulator

[5]Software available at https://github.com/brentondwalker/spark-arrivals

Appendix B: Proofs

Proof of Th. 1: We only show the proof of the sojourn time, as the proof of the waiting time follows similarly.

G|G|1 servers: We obtain from (3) for $\theta > 0$ that

$$\mathsf{E}\Big[e^{\theta T(n)}\Big] \le \sum_{\nu=1}^{n} \mathsf{E}\Big[e^{\theta S(\nu,n)}\Big]\mathsf{E}\Big[e^{-\theta A(\nu,n)}\Big],$$

where we estimated the maximum by the sum of its arguments and used the statistical independence of arrivals and service. By insertion of the (σ, ρ)-constraints from Def. 2 we have

$$\mathsf{E}\Big[e^{\theta T(n)}\Big] \le e^{\theta(\sigma_A(-\theta)+\sigma_S(\theta)+\rho_S(\theta))} \sum_{\nu=1}^{n} e^{-\theta(\rho_A(-\theta)-\rho_S(\theta))(n-\nu)}.$$

Next, we estimate

$$\begin{aligned}\sum_{\nu=1}^{n} e^{-\theta(\rho_A(-\theta)-\rho_S(\theta))(n-\nu)} &\le \sum_{\nu=0}^{\infty} \big(e^{-\theta(\rho_A(-\theta)-\rho_S(\theta))}\big)^{\nu} \\ &= \frac{1}{1-e^{-\theta(\rho_A(-\theta)-\rho_S(\theta))}},\end{aligned}$$

where we used the geometric sum for $\rho_S(\theta) < \rho_A(-\theta)$. By use of Chernoff's bound $\mathsf{P}[X \ge x] \le e^{-\theta x}\mathsf{E}\big[e^{\theta X}\big]$ we obtain

$$\mathsf{P}[T(n) \ge \tau] \le \frac{e^{\theta(\sigma_A(-\theta)+\sigma_S(\theta))}}{1-e^{-\theta(\rho_A(-\theta)-\rho_S(\theta))}} e^{\theta\rho_S(\theta)} e^{-\theta\tau}.$$

GI|GI|1 servers: From (3) we have

$$T(n) = \max_{m\in[1,n]} \{S(n-m+1,n) - A(n-m+1,n)\}.$$

For $\theta > 0$ we can write

$$\begin{aligned}&\mathsf{P}[T(n) > \tau] \\ &\quad = \mathsf{P}\left[\max_{m\in[1,n]} \left\{e^{\theta(S(n-m+1,n)-A(n-m+1,n))}\right\} > e^{\theta\tau}\right].\end{aligned}$$

Now consider the process

$$U(m) = e^{\theta(S(n-m+1,n)-A(n-m+1,n))}.$$

Using the representation of $A(m,n) = \sum_{\nu=m}^{n-1} A(\nu,\nu+1)$ and $S(m,n) = \sum_{\nu=m}^{n} S(\nu)$ by increment processes, we have

$$U(m+1) = U(m)e^{\theta(S(n-m)-A(n-m,n-m+1))}.$$

The conditional expectation can be computed as

$$\begin{aligned}&\mathsf{E}[U(m+1)|U(m),U(m-1),\dots,U(1)]\\ =&U(m)\mathsf{E}\left[e^{\theta S(n-m)}\right]\mathsf{E}\left[e^{-\theta A(n-m,n-m+1)}\right],\end{aligned}$$

where we used the independence of the inter-arrival times and the service times. If $\rho_S(\theta) \leq \rho_A(-\theta)$, it holds that $\mathsf{E}\left[e^{\theta S(n-m)}\right]\mathsf{E}\left[e^{-\theta A(n-m,n-m+1)}\right] \leq 1$ and

$$\mathsf{E}[U(m+1)|U(m),U(m-1),\dots,U(1)] \leq U(m),$$

i.e., $U(m)$ is a supermartingale. By application of Doob's inequality for submartingales [39, Theorem 3.2, p. 314] and the formulation for supermartingales [27], [44] we have for non-negative $U(m)$ for $m \geq 1$ that

$$x\mathsf{P}\left[\max_{m\in[1,n]}\{U(m)\} \geq x\right] \leq \mathsf{E}[U(1)]. \tag{39}$$

We derive

$$\mathsf{E}[U(1)] = \mathsf{E}\left[e^{\theta(S(n,n)-A(n,n))}\right] = \mathsf{E}\left[e^{\theta S(1)}\right].$$

Letting $x = e^{\theta\tau}$ we have from (39) that

$$\mathsf{P}\left[T(n) \geq \tau\right] \leq e^{\theta\rho_S(\theta)}e^{-\theta\tau},$$

which completes the proof of Th. 1. ∎

Proof of Th. 2: First, we derive the MGF of $S^{\text{net}}(m,n)$ as in Lem. 5. It follows for $\theta > 0$ that

$$\begin{aligned}\mathsf{E}\left[e^{\theta S^{\text{net}}(m,n)}\right] \leq \sum_{\nu^j: m\leq\nu^1\leq\nu^2\leq\dots\leq\nu^{h-1}\leq n} &\mathsf{E}\left[e^{\theta S^1(m,\nu^1)}\right]\\ &\mathsf{E}\left[e^{\theta S^2(\nu^1,\nu^2)}\right]\cdots\mathsf{E}\left[e^{\theta S^h(\nu^{h-1},n)}\right],\end{aligned}$$

where we estimated the maximum by the sum of its arguments and used the statistical independence of the stages. After some variable substitutions we obtain

$$\begin{aligned}\mathsf{E}\left[e^{\theta S^{\text{net}}(m,n)}\right] \leq \sum_{\nu^j\geq 0:\sum_{j=1}^{h}\nu^j=n-m} &\mathsf{E}\left[e^{\theta S^1(m,m+\nu^1)}\right]\\ \mathsf{E}\left[e^{\theta S^2(m+\nu^1,m+\nu^1+\nu^2)}\right]&\cdots\mathsf{E}\left[e^{\theta S^h(m+\sum_{j=1}^{h-1}\nu^j,m+\sum_{j=1}^{h}\nu^j)}\right].\end{aligned}$$

Given homogeneous stages that are (σ_S,ρ_S) constrained as specified by Def. 2 and using [45, Prop. 6.2] to replace the sum by a binomial coefficient, we have for $\theta > 0$ that

$$\mathsf{E}\left[e^{\theta S^{\text{net}}(m,n)}\right] \leq \binom{n-m+h-1}{h-1}e^{\theta(h\sigma_S(\theta)+\rho_S(\theta)(n-m+h))}.$$

Next, we derive for the MGF of the sojourn time from (3) for $\theta > 0$ that

$$\mathsf{E}\left[e^{\theta T(n)}\right] \leq \sum_{\nu=1}^{n}\mathsf{E}\left[e^{\theta S^{\text{net}}(\nu,n)}\right]\mathsf{E}\left[e^{-\theta A(\nu,n)}\right],$$

where we estimated the maximum by the sum of its arguments and used the statistical independence of arrivals and service. Considering (σ_A,ρ_A) constrained traffic as in Def. 2, we have

$$\begin{aligned}\mathsf{E}\left[e^{\theta T(n)}\right] \leq\ & e^{\theta(\sigma_A(-\theta)+h\sigma_S(\theta)+h\rho_S(\theta))}\\ &\sum_{\nu=1}^{n}\binom{n-\nu+h-1}{h-1}e^{-\theta(\rho_A(-\theta)-\rho_S(\theta))(n-\nu)}.\end{aligned}$$

Next, we estimate for $\rho_S(\theta) < \rho_A(-\theta)$ that

$$\begin{aligned}&\sum_{\nu=1}^{n}\binom{n-\nu+h-1}{h-1}e^{-\theta(\rho_A(-\theta)-\rho_S(\theta))(n-\nu)}\\ \leq&\sum_{\nu=0}^{\infty}\binom{\nu+h-1}{h-1}\left(e^{-\theta(\rho_A(-\theta)-\rho_S(\theta))}\right)^{\nu}\\ =&\left(1-e^{-\theta(\rho_A(-\theta)-\rho_S(\theta))}\right)^{-h},\end{aligned}$$

where we used that

$$\begin{aligned}&\sum_{\nu=0}^{\infty}\binom{\nu+h-1}{h-1}\\ &\left(e^{-\theta(\rho_A(-\theta)-\rho_S(\theta))}\right)^{\nu}\left(1-e^{-\theta(\rho_A(-\theta)-\rho_S(\theta))}\right)^{h} = 1,\end{aligned}$$

as the argument of the sum takes the form of the negative binomial probability mass function. By use of Chernoff's bound we obtain

$$\mathsf{P}[T(n) \geq \tau] \leq \frac{e^{\theta(\sigma_A(-\theta)+h\sigma_S(\theta))}}{\left(1-e^{-\theta(\rho_A(-\theta)-\rho_S(\theta))}\right)^{h}}e^{\theta h\rho_S(\theta)}e^{-\theta\tau}.$$

Finally, we insert the service parameters of the tasks $\sigma_S(\theta) = \sigma_Q(\theta) + \ln(k)/\theta$ and $\rho_S(\theta) = \rho_Q(\theta)$ from (11) and (12) for each of the fork-join stages to complete the proof. ∎

REFERENCES

[1] M. Fidler and Y. Jiang, "Non-asymptotic delay bounds for (k,l) fork-join systems and multi-stage fork-join networks," in *Proc. of IEEE INFOCOM*, Apr. 2016.

[2] J. Dean and S. Ghemawat, "Mapreduce: simplified data processing on large clusters." *Commun. ACM*, vol. 51, no. 1, pp. 107–113, 2008.

[3] M. Zaharia, M. Chowdhury, M. J. Franklin, S. Shenker, and I. Stoica, "Spark: Cluster computing with working sets," in *Proc. of USENIX Conference on Hot Topics in Cloud Computing*, 2010.

[4] F. Baccelli, A. M. Makowski, and A. Shwartz, "The fork-join queue and related systems with synchronization constraints: Stochastic ordering and computable bounds," *Adv. in Appl. Probab.*, vol. 21, no. 3, pp. 629–660, Sep. 1989.

[5] B. Kemper and M. Mandjes, "Mean sojourn times in two-queue fork-join systems: bounds and approximations," *OR Spektrum*, vol. 34, no. 3, pp. 723–742, Jul. 2012.

[6] J. Tan, X. Meng, and L. Zhang, "Delay tails in MapReduce scheduling," *ACM Sigmetrics Perf. Eval. Rev.*, vol. 40, no. 1, pp. 5–16, Jun. 2012.

[7] L. Flatto and S. Hahn, "Two parallel queues created by arrivals with two demands," *SIAM J. Appl. Math.*, vol. 44, no. 5, pp. 1041–1053, 1984.

[8] R. Nelson and A. N. Tantawi, "Approximate analysis of fork/join synchronization in parallel queues," *IEEE Trans. Comput.*, vol. 37, no. 6, pp. 739–743, Jun. 1988.

[9] A. S. Lebrecht and W. J. Knottenbelt, "Response time approximations in fork-join queues," in *Proc. of UKPEW*, Jul. 2007.

[10] S.-S. Ko and R. F. Serfozo, "Sojourn times in G/M/1 fork-join networks," *Naval Research Logistics*, vol. 55, no. 5, pp. 432–443, May 2008.

[11] X. Tan and C. Knessl, "A fork-join queueing model: Diffusion approximation, integral representations and asymptotics," *Queueing Systems*, vol. 22, no. 3-5, pp. 287–322, Sep. 1996.

[12] S. Varma and A. M. Makowski, "Interpolation approximations for symmetric fork-join queues," *Performance Evaluation*, vol. 20, no. 1-3, pp. 245–265, May 1994.

[13] E. Varki, "Mean value technique for closed fork-join networks," *ACM Sigmetrics Perf. Eval. Rev.*, vol. 27, no. 1, pp. 103–112, May 1999.

[14] F. Alomari and D. A. Menascé, "Efficient response time approximations for multiclass fork and join queues in open and closed queuing networks," *IEEE Trans. Parallel Distrib. Syst.*, vol. 25, no. 6, pp. 1437–1446, Jun. 2014.

[15] A. Rizk, F. Poloczek, and F. Ciucu, "Stochastic bounds in fork-join queueing systems under full and partial mapping," *Queueing Systems: Theory and Applications*, vol. 83, no. 3, pp. 261–291, Aug. 2016.

[16] G. Kesidis, B. Urgaonkar, Y. Shan, S. Kamarava, and J. Liebeherr, "Network calculus for parallel processing," in *Proc. of MAMA Workshop at ACM SIGMETRICS*, Jun. 2015.

[17] Y. Han and A. M. Makowski, "Resequencing delays under multipath routing – asymptotics in a simple queueing model," in *Proc. of IEEE INFOCOM*, Apr. 2006.

[18] Y. Xia and D. Tse, "On the large deviation of resequencing queue size: 2-M/M/1 case," *IEEE Trans. Inf. Theory*, vol. 54, no. 9, pp. 4107–4118, Sep. 2008.

[19] Y. Gao and Y. Q. Zhao, "Large deviations for re-sequencing buffer size," *IEEE Trans. Inf. Theory*, vol. 58, no. 2, pp. 1003–1009, Feb. 2012.

[20] P. Harrison and S. Zertal, "Queueing models with maxima of service times," in *Proc. of TOOLS*, Sep. 2003, pp. 152–168.

[21] G. Joshi, Y. Liu, and E. Soljanin, "On the delay-storage trade-off in content download from coded distributed storage systems," *IEEE J. Sel. Areas Commun.*, vol. 32, no. 5, pp. 989–997, May 2014.

[22] F. Poloczek and F. Ciucu, "Contrasting effects of replication in parallel systems: From overload to underload and back," arXiv, Tech. Rep. arXiv:1602.07978v1, Feb. 2016.

[23] Q. Yin, Y. Jiang, S. Jiang, and P. Y. Kong, "Analysis of generalized stochastically bounded bursty traffic for communication networks," in *Proc. of IEEE LCN*, Nov. 2002, pp. 141–149.

[24] C.-S. Chang, *Performance Guarantees in Communication Networks*. Springer-Verlag, 2000.

[25] F. Baccelli, G. Cohen, G. J. Olsder, and J.-P. Quadrat, *Synchronization and Linearity: An Algebra for Discrete Event Systems*. Wiley, 1992.

[26] J. Xie and Y. Jiang, "Stochastic network calculus models under max-plus algebra," in *Proc. of IEEE GLOBECOM*, 2009.

[27] Y. Jiang, "Network calculus and queueing theory: Two sides of one coin," in *Proc. of VALUETOOLS*, 2009, pp. 1–12, Invited Paper.

[28] R. Lübben, M. Fidler, and J. Liebeherr, "Stochastic bandwidth estimation in networks with random service," *IEEE/ACM Trans. Netw.*, vol. 22, no. 2, pp. 484–497, Apr. 2014.

[29] R. L. Cruz, "A calculus for network delay, part I and II: Network elements in isolation and network analysis," *IEEE Trans. Inf. Theory*, vol. 37, no. 1, pp. 114–141, Jan. 1991.

[30] J.-Y. Le Boudec and P. Thiran, *Network Calculus A Theory of Deterministic Queuing Systems for the Internet*. Springer-Verlag, 2001.

[31] A. Burchard, J. Liebeherr, and S. Patek, "A min-plus calculus for end-to-end statistical service guarantees," *IEEE Trans. Inf. Theory*, vol. 52, no. 9, pp. 4105–4114, Aug. 2006.

[32] F. Ciucu, A. Burchard, and J. Liebeherr, "Scaling properties of statistical end-to-end bounds in the network calculus," *IEEE/ACM Trans. Netw.*, vol. 14, no. 6, pp. 2300–2312, Jun. 2006.

[33] M. Fidler, "An end-to-end probabilistic network calculus with moment generating functions," in *Proc. of IWQoS*, Jun. 2006, pp. 261–270.

[34] Y. Jiang and Y. Liu, *Stochastic Network Calculus*. Springer-Verlag, Sep. 2008.

[35] M. Fidler, "A survey of deterministic and stochastic service curve models in the network calculus," *IEEE Commun. Surveys Tuts.*, vol. 12, no. 1, pp. 59–86, 2010.

[36] F. Ciucu and J. Schmitt, "Perspectives on network calculus - no free lunch but still good value," in *Proc. of ACM SIGCOMM*, Aug. 2012, pp. 311–322.

[37] M. Fidler and A. Rizk, "A guide to the stochastic network calculus," *IEEE Commun. Surveys Tuts.*, vol. 17, no. 1, pp. 92–105, 2015.

[38] J. F. C. Kingman, "A martingale inequality in the theory of queues," *Math. Proc. Cambridge*, vol. 60, no. 2, pp. 359–361, Apr. 1964.

[39] J. L. Doob, *Stochastic Processes*. Wiley, 1953.

[40] I. Adan and J. Resing, *Queueing Systems*. Eindhoven University of Technology, Department of Mathematics and Computing Science, 2015.

[41] Y. Jiang, "A basic stochastic network calculus," in *Proc. of ACM SIGCOMM*, Oct. 2006, pp. 123–134.

[42] D. Gross, J. F. Shortle, J. M. Thompson, and C. M. Harris, *Fundamentals of Queueing Theory*, 4th ed. Wiley, 2008.

[43] J.-Y. Le Boudec, *Performance evaluation of computer and communication systems*, ser. Computer and communication sciences. EPFL Press London, 2010.

[44] Y. Jiang, "A note on applying stochastic network calculus," Tech. Rep., 2010.

[45] S. Ross, *A First Course in Probability*, 7th ed. Pearson Prentice Hall, 2006.